2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

RICE: Reactive Interaction Controller for Cluttered Canopy Environment

Nidhi Homey Parayil[1], Thierry Peynot[2], and Chris Lehnert[3]

***Abstract*— Robotic motion in dense, cluttered environments such as agricultural canopies presents significant challenges due to physical and visual occlusion caused by leaves and branches. Traditional vision-based or model-dependent approaches often fail in these settings, where physical interaction without damaging foliage and branches is necessary to reach a target. We present a novel reactive controller that enables safe navigation for a robotic arm in a contact-rich, cluttered, deformable environment using end-effector position and real-time tactile feedback. Our proposed framework's interaction strategy is based on a trade-off between minimizing disturbance by maneuvering around obstacles and pushing through them to move towards the target. We show that over 35 trials in 3 experimental plant setups with an occluded target, the proposed controller successfully reached the target in all trials without breaking any branches and outperformed the established control strategy for dense foliage in reliability and adaptability. This work lays the foundation for safe, adaptive interaction in cluttered, contact-rich deformable environments, enabling future agricultural tasks such as pruning and harvesting in plant canopies.**

I. Introduction

Robots struggle to operate in an agricultural environment due to dense and unstructured clutter, such as overlapping leaves and branches [1]. This clutter creates both physical obstructions, which require robots to interact with or navigate around obstacles, and visual occlusions, which hinder perception and path planning toward targets like fruits. When navigating cluttered environments, there are generally three possible strategies: pushing through obstacles, navigating around them, or adaptively combining both [2]. However, most robotic systems are limited to the first two options [3], and obstacle avoidance is the most common approach in agricultural applications [4]. As a result, recent developments in agricultural automation emphasize the selection of obstacle-free paths for tasks such as harvesting or pruning to minimize interactions with clutter, thus limiting the application of robotic manipulators in agriculture [5]. This paper addresses the challenge of physical interaction and navigation to reach a given target in cluttered agricultural settings effectively.

The primary challenge for the deployment of robots within plant environments is generating motion that avoids damaging delicate structures like stems and branches [4]. This demands control strategies that regulate both trajectory and applied force, often guided by visual or spatial information [6]. In a cluttered environment, occlusions hinder perception, making it difficult to generate reliable trajectories. Hybrid controllers are often used for physical interaction by switching between position and force control based on task-specific criteria [7]. However, clutter often causes frequent physical contact, requiring *simultaneous* control of the path and force profiles [6]. During contact, the robot's dynamics become coupled with those of the environment, complicating control and increasing reliance on accurate environmental models.

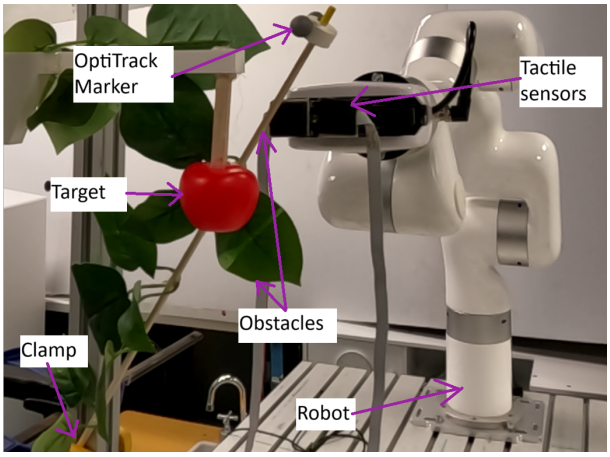

Fig. 1. Quantitative experimental setup: In a cluttered canopy, the robot navigates toward a given target, encountering branches and leaves as obstacles. OptiTrack markers measure branch disturbance. The robot adaptively pushes through or moves around obstacles to minimize disturbance using tactile feedback.

How do people handle visual occlusion when harvesting fruit among leaves and branches? While vision provides valuable spatial cues, it is often incomplete in dense foliage, where targets are partially or fully obscured. In such cases, humans naturally rely on tactile feedback to sense resistance. Leaves and branches are deformable, i.e. their shape and position change under contact, making the environment highly dynamic and difficult to model precisely. As we move, we feel how leaves and branches shift or bend, adjusting our motion to avoid excessive force or plant damage. In essence, we perform tactile servoing [8]. Without accurate dynamic models, it is hard to predict how the environment will respond to interaction. To handle such rapid and unpredictable changes, a controller must react in real time. Since classical planners struggle to accommodate fast, dynamic interactions [9], we use a reactive controller for this work.

To address the challenges posed by cluttered canopies, we propose a controller architecture called the Reactive

All authors are with the QUT Centre of Robotics, Queensland University of Technology (QUT), Brisbane QLD 4001, Australia (e-mail: nidhihomey.parayil@hdr.qut.edu.au). N.H.P., T.P., and C.L. acknowledge continued support from QUT through the Centre for Robotics. We wish to acknowledge the support of the Research Engineering Facility (REF) team at QUT, particularly Alec Tutin and Peter Smith, for the provision of expertise and research infrastructure in enabling this project.

Interaction Controller for Cluttered Canopy Environments (RICE). The key contributions of this paper are as follows. We introduce a hierarchical, model-free control architecture that enables safe navigation in a contact-rich, cluttered, deformable environment using end-effector position and real-time tactile feedback to reach a target without damaging foliage and branches. We propose an optimal motion strategy that adaptively balances between pushing through to reach the target and maneuvering around obstacles to minimize interaction forces. To support quantitative analysis, we design a plant deformation measurement setup using a motion capture system, enabling capture of environmental interactions. Finally, we validate our approach quantitatively in custom-built, trackable mock plant environments across 30 distinct trials, as well as qualitatively in denser foliage with 5 trials, demonstrating superior robustness and adaptability compared to existing control strategies for agricultural environments.

II. Prior work

Navigating in a cluttered environment can involve various strategies depending on the task, but a common goal in agricultural tasks is to minimize damage to the robot or the environment. This includes avoiding damage to stems, leaves, and trellis wires. Since free space is limited, the robot must often push through the environment to create a path. Motion planners designed for physical interaction [10]–[12] are typically slow because they incorporate kino-dynamic constraints. As a result, they fail to react quickly to unexpected obstacles in cluttered environments. Reactive controllers make short horizon decisions based on real-time environment interaction feedback, enabling adaptation to changing dynamics in unpredictable environments [13]. These systems often combine two control layers running at different frequencies [14]: a high-level controller that responds to environmental changes and a low-level controller that manages the robot's dynamics. Haviland *et al.* [15] developed a reactive controller that uses quadratic programming on the robot model to optimize manipulability at the high level and employs Resolved Rate Motion Control (RRMC) at the low level. Their controller achieves faster performance than classical planners, but it does not support physical interaction because it does not use the model of the environment.

Interaction control in cluttered environments involves both contact and non-contact phases, requiring the controller to regulate both position and force depending on the task context. Force-based controllers like admittance and impedance maintain desired force but only function when in contact [6]. Position-based controllers can push obstacles but focus on trajectory, instead of force. Hybrid controllers use position control in non-contact situations, and switch to force control once contact begins or when forces exceed a threshold [16]. Tasks, like surface cleaning, apply force control along one axis and position along another [7]. Iskandar *et al.* developed a method to switch between force and position control independently on each axis [7], [17]. Wang *et. al* used an approximate model of environment dynamics combined with a fuzzy logic-based hybrid controller to handle uncertainty [18]. Building such models is difficult in cluttered, deformable environments, limiting their use in real-world agricultural settings.

Model Predictive Controllers (MPCs) are popular for interaction tasks because they can integrate multiple objectives and constraints effectively [19]. However, this approach requires a model of the environment to define its constraints. In an agricultural environment, controlling interaction force precisely is often impractical, as a force suitable for one branch may damage another. One study in agriculture combined MPC with skin force sensors to navigate clutter made of leaves and logs [20]. The method demonstrated the robot pushing through compliant objects like leaves while avoiding rigid ones like logs. It does not extend to interactions with elements like branches, which present additional challenges due to their deformability and fragility, and may require the robot to selectively push or maneuver around them.

An alternative approach in agriculture for navigating clutter involves manipulating the environment. One method uses a zig-zag pattern with multi-directional pushes to break contact with unripe strawberries, allowing access to hidden ripe ones [21]. This strategy, though effective for fruit, struggles with leaves and stems, which have different dynamic properties. Lehnert et al. proposed moving the robot to improve visibility when occlusion occurs in cluttered environments [22]. While this addresses visual occlusion, it does not solve the challenge of physical obstruction. Other methods plan around uncertainty by choosing paths with fewer obstacles or rearranging objects [2], [23]. For crop monitoring, robotic arms have also been used to gently move leaves [24] or branches [25] aside to reveal fruit. Although manipulating obstacles can reduce occlusion, clearing all foliage can be time-consuming.

Learning-based control techniques are gaining popularity for robots interacting with deformable objects. Most work focuses on 2D deformable objects such as cloth and ropes [26], [27]. Some approaches have addressed table-top clutter, such as a robot reaching a target hidden within soft objects like foam balls [3]. These environments are typically constrained to fixed surfaces with limited depth variation, and interaction occurs primarily in 2D. Agricultural environments are less structured and involve more complex, three-dimensional interactions with deformable elements like leaves and branches. Medical robotics has explored complex, unstructured, 3D deformable environments. For example, researchers have guided needles through soft tissue using demonstrations, inverse reinforcement learning, and accurate simulation models built from patient scans [28]. While these techniques are accurate, they rely on extensive, high-fidelity data such as CT/MRI scans and anatomical models. Agricultural environments lack such data availability. Plant geometry and material properties vary widely across species, growth stages, and even environmental conditions. Although some work has explored plant simulation [29], current tools and datasets remain insufficient to support detailed models for learning-based models in agriculture.

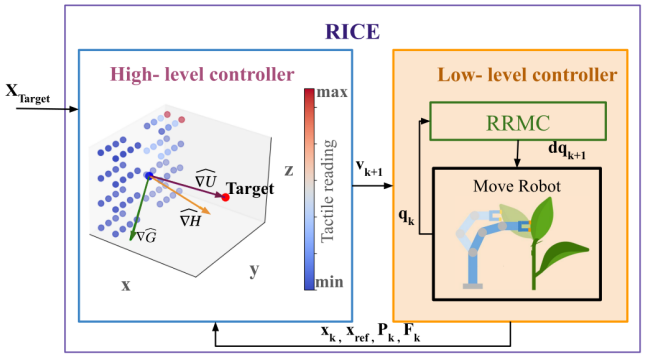

Fig. 2. Overall controller architecture showing a hierarchical scheme. The high-level controller uses end-effector position ($\mathbf{x}_k$), tactile forces ($\mathbf{F}_k$), and taxel positions ($\mathbf{P}_k$) to compute a velocity command toward the target ($\mathbf{x}_{Target}$). The low-level controller is a Resolved Rate Motion Controller (RRMC), which computes joint velocities from desired end-effector velocities using joint-state feedback.

Moving a robot through cluttered agricultural environments with physical interaction calls for a model-free, reactive controller capable of regulating position and force trajectory. The controller must make real-time decisions based on local tactile feedback to minimise environmental damage while effectively maneuvering around deformable plant structures. We propose a reactive control method that combines end-effector position with tactile force feedback, enabling a strategy that balances pushing through to reach the target and maneuvering around obstacles to minimize disturbance in dense clutter.

III. RICE: Reactive Interaction Controller for Cluttered Canopy Environment

In cluttered agricultural environments, humans instinctively move towards the target by choosing a path that optimally balances distance and environmental resistance. We move around stems, which are narrow and rigid, and push through compliant leaves. These conditions require a continuous trade-off between pushing through and moving around obstacles, rather than relying on a single optimal strategy. We exploit the fact that this trade-off is correlated with the spatial coverage of each object on the tactile sensors. Trellis wires typically cover minimal area with most resistance, branches cover larger areas but can deform, and leaves may cover the entire sensor array yet allow compliant interaction. Drawing on this insight, we design our objective function to:

- prioritize moving directly towards the objects when the robot detects no contact or uniform, low-resistance contact (e.g., from leaves),
- maneuver around partial obstructions (e.g., stems) where excessive force could lead to plant damage.

We do not address interactions with large branches that completely occlude the sensor surface, as we consider these cases beyond the scope of this work. Our focus is to reach a target for tasks such as harvesting or pruning tasks involving plants like tomatoes, berry bushes and grapevines, where such large obstructions are uncommon.

A. Sensor configuration

To sense interaction forces in cluttered environments, our system uses two tactile arrays mounted on each gripper finger (Fig. 3(a)). Each array has an $n \times n$ taxel grid, each taxel producing a three-axis force measurements per low-level control step. The end effector frame EE is defined as $\mathbf{x}_{\text{EE}}, \mathbf{y}_{\text{EE}}, \mathbf{z}_{\text{EE}}$ for forward, lateral, and vertical directions respective to the gripper. Given planar placement and known offsets in $\mathbf{y_{EE}}$ and $\mathbf{z_{EE}}$, the system captures $2n^2$ measurements in the end effector frame at each cycle. To obtain force variation along the $\mathbf{x_{EE}}$ axis, we record both tactile values and taxel positions across consecutive time steps of the low-level controller within a single high-level control cycle, as shown in Fig. 3(b). Although full skin-like coverage would be ideal, sensing is limited to the gripper tips due to hardware and cost constraints.

B. Controller Architecture

RICE is a hierarchical control architecture comprising a high-level, optimization-based controller operating at frequency f_{H} and a low-level controller running at a higher frequency f_{L}. Their respective time steps are $T = 1/f_{\text{H}}$ and $\tau = 1/f_{\text{L}}$, and $k \in \mathbb{N}$ is used for indexing high-level steps. The frequency ratio $j = f_{\text{L}}/f_{\text{H}} \in \mathbb{N}$ indicates that the low-level controller executes j times between $k-1$ and k. At each high-level step k, the system collects j tactile frames from the low-level controller, denoted as $\{\mathbf{F}_{k_1}, \ldots, \mathbf{F}_{k_j}\}$. Each frame $\mathbf{F}_{k_m} \in \mathbb{R}^{N\times 3}$ contains 3D force vectors $\mathbf{f} = [f_x, f_y, f_z] \in \mathbb{R}^3$ from $N = 2n^2$ taxels:

$$\mathbf{F}_{k_m} = \begin{bmatrix} f_{x1} & f_{y1} & f_{z1} \\ & \vdots & \\ f_{xN} & f_{yN} & f_{zN} \end{bmatrix}, \quad m = 1, \ldots, j$$

To aggregate tactile data over the high-level interval, these frames are concatenated horizontally $\mathbf{F}_k = \begin{bmatrix} \mathbf{F}_{k_1} & \cdots & \mathbf{F}_{k_s} \end{bmatrix}^\top \in \mathbb{R}^{s\times 3}$, where $s = N \cdot j$. Corresponding taxel positions are represented by $\mathbf{P}_k = [\mathbf{P}_{k_1} \cdots \mathbf{P}_{k_s}]^\top \in \mathbb{R}^{s\times 3}$, with each column $\mathbf{p}_{ki} = [p_{kix}, p_{kiy}, p_{kiy}] \in \mathbb{R}^3$ indicating the spatial coordinates of the respective taxel. The high-level controller utilizes the end-effector position $\mathbf{x}_k \in \mathbb{R}^3$, reference end effector position (at $\tau = 0$ time frame) $\mathbf{x}_{ref} \in \mathbb{R}^3$, target position $\mathbf{x}_{Target} \in \mathbb{R}^3$, taxel positions $\mathbf{P}_k$, and tactile feedback $\mathbf{F}_k$ to compute an optimal end-effector velocity toward the target. The low-level controller translates this velocity into joint velocity $\dot{\mathbf{q}}_k \in \mathbb{R}^b$ commands using joint position $\mathbf{q}_k \in \mathbb{R}^b$ feedback where b is the degree of freedom of the robot.

C. High-level controller

RICE computes the optimal velocity vector $\mathbf{v}_{k+1} \in \mathbb{R}^3$ at each high-level time step $T = k$ by balancing two objectives: reaching the target and minimizing interaction forces. The objective function $H_k \in \mathbb{R}^s$ is defined as a weighted combination of the target reach cost $U(\mathbf{x}_k, \mathbf{x}_{\text{Target}}) \in \mathbb{R}$ and interaction force cost $G(\mathbf{x}_k, \mathbf{P}_k, \mathbf{F}_k, \mathbf{x}_{\text{ref}}) \in \mathbb{R}^s$, scaled by

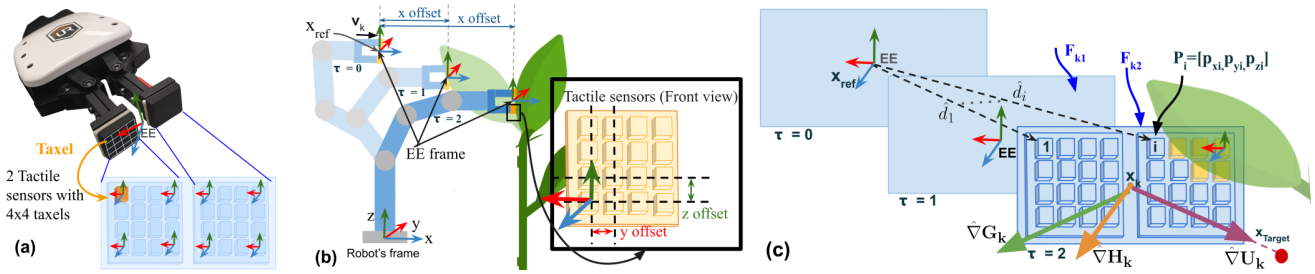

Fig. 3. Left: Each fingertip is equipped with a 4×4 tactile sensor array ($n = 4$), providing 16 taxels per sensor. EE represents the end effector's reference frame. Centre: Side view of the robot interacting with a leaf during one high-level control step, with the low-level controller running from $\tau = 0$ to $\tau = 2$ ($j = 2$). Forces along $\mathbf{x_{EE}}$ and $\mathbf{z_{EE}}$ are measured at each τ, while $\mathbf{x_{EE}}$ axis variation is estimated across steps. Right: The tactile sensors respond to contact with the leaf (yellow taxels). RICE optimization computes motion direction $\hat{\nabla} H$ as a weighted sum of the target reach cost $\hat{\nabla} U$ and interaction force cost $\hat{\nabla} G$. $\hat{d}_i$ is the direction vector from reference position to i^{th} taxel and $\mathbf{p_i}$ is the corresponding taxel position.

scalar hyperparameters $w_x, w_f \in \mathbb{R}^+$:

$$H_k = w_x U(\mathbf{x}_k, \mathbf{x}_{\text{Target}}) + w_f G(\mathbf{x}_k, \mathbf{P}_k, \mathbf{F}_k, \mathbf{x}_{\text{ref}}). \quad (1)$$

The optimal motion is given by the gradient descent of the objective function, and its gradient is given by:

$$\nabla \mathbf{H_k} = w_x \, \hat{\nabla} \mathbf{U}(\mathbf{x_k}, \mathbf{x}_{\text{Target}}) + w_f \, \hat{\nabla} \mathbf{G}(\mathbf{x_k}, \mathbf{P_k}, \mathbf{F_k}, \mathbf{x}_{\text{ref}}), \quad (2)$$

where $\hat{\nabla}$ denotes a normalized gradient to account for differences in units and scale. Here, w_x emphasizes progress toward the goal (especially in free space), while w_f promotes deviation to reduce contact forces (e.g., when interacting with stems). This formulation allows the controller to adapt its trajectory in real time, balancing goal-directed motion with tactile-aware compliance (Fig. 3(c)).

1) Target reach cost: The gradient of the target reach cost $\nabla \mathbf{U_k}$ is computed from a potential function inspired by potential field method [30], where the robot is attracted toward a goal. The cost defined as the squared Euclidean distance between the current position $\mathbf{x}_k$ and the target $\mathbf{x}_{\text{Target}}$ given by $U_k = \|\mathbf{x}_{\text{Target}} - \mathbf{x}_k\|^2$. The corresponding gradient is$\nabla \mathbf{U_k} = -2(\mathbf{x}_{\text{Target}} - \mathbf{x}_k)$. The normalised gradient is then:

$$\hat{\nabla} \mathbf{U_k} = \frac{\nabla \mathbf{U_k}}{\|\nabla \mathbf{U_k}\|}. \quad (3)$$

2) Interaction force cost: The interaction force cost G is computed using tactile feedback $\mathbf{F}_k$, corresponding taxel positions $\mathbf{P}_k$, reference end-effector position $\mathbf{x}_{\text{ref}}$, and current end-effector position $\mathbf{x}_k$ during the kth step of the high-level controller. The idea of using an interaction force cost gradient across all axes to guide motion, minimizing contact forces in clutter, is inspired by the approach of Lehnert *et al.* [22], where visual gradients were used to move towards regions of higher visibility of fruits. We define direction vectors from the reference point $\mathbf{x}_{\text{ref}}$ to each taxel position as:

$$\hat{\mathbf{D}_k} = \begin{bmatrix} \hat{\mathbf{d}_{k1}} & \cdots & \hat{\mathbf{d}_{ks}} \end{bmatrix}^\top, \quad \hat{\mathbf{d}_{ki}} = \frac{\mathbf{p}_{ki} - \mathbf{x}_{\text{ref}}}{\|\mathbf{p}_{ki} - \mathbf{x}_{\text{ref}}\|}, \quad (4)$$

where $i \in [0, s]$, $\mathbf{p}_{ki} \in \mathbb{R}^3$ denotes the Cartesian coordinates of the ith taxel, and $\hat{\mathbf{D}_k} \in \mathbb{R}^{s \times 3}$ contains unit vectors pointing from $\mathbf{x}_{\text{ref}}$ to each taxel. The force magnitude vector $\mathbf{G}_k \in \mathbb{R}^s$ is computed as the Euclidean norm of each taxel's 3D force:

$$G_k = \begin{bmatrix} \|\mathbf{f}_1\|_2 \cdots \|\mathbf{f}_s\|_2 \end{bmatrix}^\top. \quad (5)$$

The reference magnitude is given by $g_{\text{ref}} = \|\mathbf{f}_{\text{ref}}\|_2$, where $\mathbf{f}_{\text{ref}}$ is the average tactile force vector at $\tau = 0$. We define the interaction force cost deviation as:

$$\Delta \mathbf{G_k} = \begin{bmatrix} \|\mathbf{f}_1\|_2 - \|\mathbf{f}_{\text{ref}}\|_2 \\ \vdots \\ \|\mathbf{f}_s\|_2 - \|\mathbf{f}_{\text{ref}}\|_2 \end{bmatrix}. \quad (6)$$

Using the directional derivative approximation $\nabla \mathbf{G} \hat{\mathbf{d}}_i \approx \frac{g_i - g_{\text{ref}}}{\|\mathbf{d}_i\|}$, we model the deviation as, $\Delta \mathbf{G_k} = \hat{\mathbf{D}}_k \nabla \mathbf{G_k}$. The spatial gradient $\nabla \mathbf{G}_\mathbf{k}^* \in \mathbb{R}^3$ is then estimated using a least-squares solution:

$$\nabla \mathbf{G}_\mathbf{k}^* = (\hat{\mathbf{D}}_k^\top \hat{\mathbf{D}}_k)^{-1} \hat{\mathbf{D}}_k^\top \Delta \mathbf{G_k} \quad (7)$$

and the normalized gradient is given by $\hat{\nabla} \mathbf{G_k}$.

3) Optimal Velocity Calculation: The optimal velocity vector $\mathbf{v}_{k+1} \in \mathbb{R}^3$ is computed using a normalised gradient descent method, based on the gradient vector $\nabla \mathbf{H}_k$ and task-dependent step size $\alpha \in \mathbb{R}$. The velocity is given by Eq. 8:

$$\mathbf{v}_{k+1} = -\alpha \cdot \frac{\nabla \mathbf{H}_k}{|\nabla \mathbf{H}_k|}. \quad (8)$$

D. Low-level controller

We used Resolved-Rate Motion Controller (RRMC) [31] for low-level motion control. The low-level controller is integrated with tactile sensing. RRMC maps the optimal end effector velocity $\mathbf{v}_{k+1}$ calculated by the high-level controller into joint velocities $\dot{\mathbf{q}}_k$ using the current joint configuration $\mathbf{q}_k$ (see Fig 2) using Eq. 9

$$\dot{\mathbf{q}}_{k+1} = J(\mathbf{q}_k)^+ \mathbf{v}_{k+1}. \quad (9)$$

RRMC uses the pseudoinverse of the Jacobian matrix J^+, offering improved robustness near singularities compared to traditional inverse kinematics approaches [15]. The high-level controller operates reactively, updating its velocity command based on real-time sensor feedback while allowing sufficient time for the low-level controller to execute each motion segment.

IV. System and Evaluation Framework

We evaluated our controller against baseline and state-of-the-art methods using two custom trackable mock plant setups (Fig. 1) and one unstructured artificial plant. This section outlines the robotic system, mock plant environment, evaluation metrics, and testing procedure.

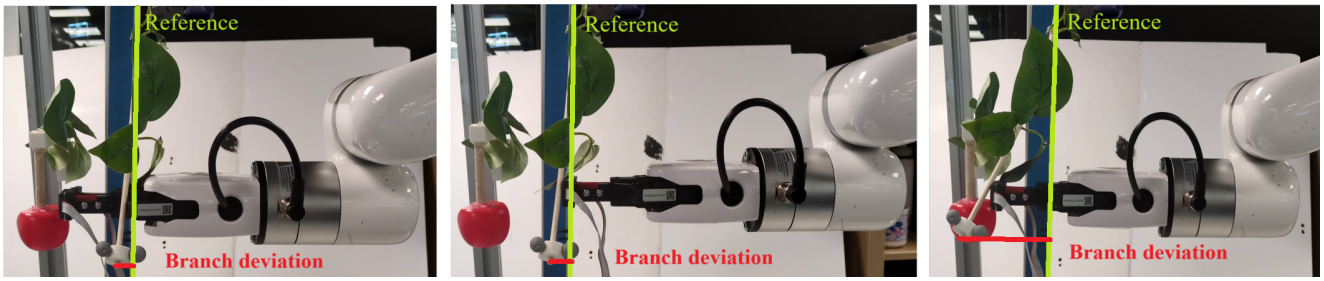

Fig. 4. Images comparing end states for the three controllers interacting with a single-branch experiment. RICE (left image) moving around the obstacle to reach the target; the hybrid controller (centre image) pushing until the force threshold is exceeded, failing to reach the target; and the position controller (right image) pushing through the branch to reach the target.

TABLE I

EXPERIMENTAL PARAMETERS USED IN EVALUATION

Parameter	Value
Balsa wood cross-section	circular ∅(5/10/12)mm, square (5 × 5)mm
Number of initial states	5
Obstacle positions	10
Branch orientations (deg)	[60, 30, 0, -30, -60]
Step size (speed)	0.01 m/s
Target positions	15

1) Robotic System: We use a UFactory XArm6 robot with a compact gripper [32] and two Xela uSPa44 tactile sensors [33] mounted via a custom 3D-printed fixture (Fig.1). Each sensor has a 4×4 taxel grid. The high-level controller runs at 50 Hz, the low-level at 100 Hz, and the tactile data at 120 Hz. The optimization algorithm averages 1.7 ms per step. The robot stops if it comes to a halt due to contact with a branch, large path deviation, or collisions with objects outside branches/leaves. Experiments run on Ubuntu 20.04 with an NVIDIA RTX 3090 GPU. OptiTrack cameras track branch deviations [34].

2) Custom trackable mock plant setup: Trackable mock plants were built using balsa wood branches with artificial leaves, chosen for their realistic deformation and breakage properties. To reflect the circular cross-section of natural branches, we tested three different circular diameters, and to assess sensitivity to shape we included one square branch. All branches were securely mounted to frames to simulate attachment to a larger plant. OptiTrack markers were placed at each branch tip for measuring deviation (Fig. 1). We tested four branch types with five initial joint configurations, ten obstacle placements, and fifteen target positions to span various contact scenarios and approach angles.Varying the branch orientation and position also resulted in different branch lengths, which in turn affected the stiffness and deformability. To ensure contact at the gripper tip where sensors are mounted, all branches were aligned parallel to it. Target positions were chosen to ensure reachability. Each trial recorded joint states and end-effector motion. Table I lists all the parameter variations used for our experiments.

3) Comparison: To our knowledge, no existing controller is designed for a safe motion in contact-rich, cluttered environments (e.g., plant canopies) where the robot must gently push through foliage with limited perception and no environmental model. The desired controller in such a setting should have two objectives: i) reach the goal and ii) minimize the disturbance. Currently, no single controller in the literature achieves both objectives simultaneously in agricultural settings. Therefore, we compare RICE against two established controllers in the literature, each representing one of these objectives:

(i) *Position Controller* (adapted from [21]): drives the end effector directly toward the target, effectively pushing through obstacles and ignoring contact forces, a simple approach suited when there is less risk of damage.

(ii) *Hybrid Controller* (adapted from [35]): You *et al.* applied force control in the y and z axes to guide branches into their curved blade, but our setup differs as we use a planar end-effector designed to reach into foliage, with branches typically in front of the robot. Therefore, we apply admittance control along the forward x-axis (the primary contact direction) and retain position control elsewhere. We set the desired contact force to $1\,\mathrm{N}$ for safe interaction with thin branches, with admittance gains $\mathrm{diag}(M) = [0, 0, 0, 100, 0, 0]$ and $\mathrm{diag}(B) = [0, 0, 0, 50, 0, 0]$. These values were determined experimentally for optimal performance in our setup. Unlike You et al., who relied on vision, our controller is designed for cluttered environments with visual occlusion; thus, we use proprioceptive feedback (external force and end effector position) and a provided target pose instead.

All controllers use resolved rate motion control (RRMC) for low-level execution with a target speed of $0.01\,\mathrm{m/s}$ to ensure consistent evaluation. These implementations ensure that each baseline was fairly tuned and evaluated, providing a meaningful comparison to RICE. We cannot compare with MPC, which would require an accurate plant model (not available in the literature), or learning-based controllers (imitation / RL), which require substantial task-specific data or high-fidelity plant simulations (Sec. II).

4) Evaluation Metrics: To compare our controller with baseline methods, we evaluated performance using both quantitative and qualitative metrics. Quantitative metrics were used to measure branch disturbance and target deviation, while also classifying overall trial success. Qualitative evaluation was done visually to assess behavior in dense clutter, where tracking individual branches was not feasible. For quantitative analysis, we used simpler setups without excessive clutter to accurately track branch motion. *Branch deviation:* Defined as the maximum deviation of each

branch from its initial position to measure environmental disturbance. For multi-branch setups, total disturbance is the sum of individual deviations. Deviation from target: Distance between the robot's end-effector and the target at the end of the trial. *No-Break Reach Rate:* Binary classification of success, defined as the percentage of trials where the robot reached the target without breaking any branches.

5) Experimental Setups: For quantitative tests, we used a trackable mock plant setup, and for qualitative visual tests, an artificial plant. We conducted (A) *Parameter sweep*. Varied the force weight $w_f \in [0.2, 3]$ to select the optimal value, keeping other parameters constant and the branch positioned vertically in front of the sensor. Following this, four experiments were conducted: (B) *Single-branch:* Compared controllers using 20 trials per controller with varying combinations of the parameters listed in Table I (C) *Multi-branch:* Evaluated performance with two branches using 10 trials per controller, varying parameters from Table I. (D) *RICE Reliability Evaluation:* Assessed RICE's reliability with repeated runs, totaling 100 single-branch and 25 two-branch trials. (E) *Artificial plant (qualitative):* five tests to visually assess behavior when motion-capture tracking is unreliable due to occlusions.

V. Results

Results and analysis for each experiment are detailed here.

A. Parameter Sweep

Maximum branch deviation across different force weight values (w_f) is shown in Fig. 5. Since both the target cost and force cost gradients are normalized ($\hat{\nabla}U_k$,$\hat{\nabla}G_k$), tuning only w_f is sufficient to balance goal-seeking and contact avoidance behavior. At low values ($w_f < 1.2$), the controller tended to follow a straight-line trajectory, often pushing through obstacles and causing significant branch disturbance. In contrast, high values ($w_f > 2$) led to repeated ineffective behavior: The robot would step back after contact but re-attempt the same trajectory, resulting in contact loops. Intermediate values ($w_f = 1.4$ to 1.8) reduced disturbance but did not provide sufficient backward motion for effective recovery from sudden contacts. The best overall performance was achieved at $w_f = 2$, where the robot exhibited stable and adaptive behavior. Upon contact, it made small backward and lateral adjustments, allowing branches to recover from the disturbance and enabling it to re-approach the target from a different angle. This strategy reduced repeated contact and prevented motion loops critical in deformable, cluttered environments where prolonged interaction increases the risk of damage. Based on these findings, $w_f = 2$ was used for all subsequent experiments. Note that this value is specific to the current setup and may require retuning for different plant geometries or tasks.

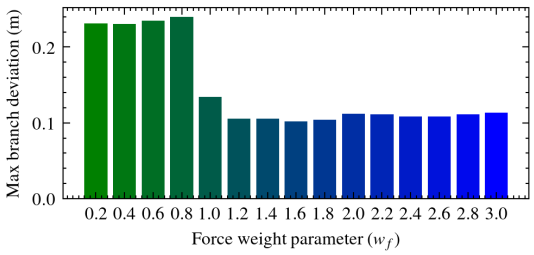

Fig. 5. Maximum branch deviation for different force weight parameters. The robot exhibits motion around an obstacle for a force weight larger than 1.2. The variation in deviation of w_f from 1.2 to 3 is due to the dynamic, uncertain nature of the branch.

B. Single Branch

Across 20 trials, the RICE controller successfully reached the target in every case without breaking any branches (Table II). It achieved a median branch deviation of 22mm, adaptively moving through foliage by pushing gently against leaf edges while avoiding stiffer midribs. The position-based controller also reached the target in all trials but caused the most disturbance, with five branch breaks and a median branch deviation of 52mm. Following a fixed path, it often forced through obstacles, even breaking thicker branches that did not deform easily. The hybrid controller succeeded in 9 of 20 trials, with three branch breaks. Its performance depended on branch orientation and target direction. It exhibited four outcomes: (i) reached the target when the position controller on the other axis guided motion around obstacles, (ii) got stuck after reaching the 1 N contact limit, (iii) deviated significantly when moving along both axes, or (iv) broke a branch when it was too rigid typically when branches were clamped closer to the edge, making them stiff. Median branch deviation was 33 mm, and median target deviation was 31.1 mm. The results across all trials for the three controllers are summarized in Fig. 6 (Centre), highlighting RICE's consistent ability to minimize branch disturbance. An example trajectory from a single test is shown in Fig.6 (Left), where the position controller pushed through the branch, the hybrid controller stopped upon contact, and RICE maneuvered around it. The corresponding branch motion, recorded using the OptiTrack system, is shown in Fig.7, indicating that RICE caused the least disturbance in this particular trial. The final end-effector and branch positions for this same test are illustrated in Fig. 4.

C. Multi-Branch

The RICE controller reached the target in all ten trials, with total branch deviation of 35mm (refer to Fig. 6 Centre) and no breakage, resulting in a 100% No-Break Reach rate (see Table II). It moved around obstacles without causing damage, showing strong performance in complex, multi-branch setups. The hybrid controller reached the target in only one trial and failed to move past the first branch in the rest, resulting in a 10% No-Break Reach rate. It had a median branch movement of 21.5mm and a deviation from target of 63mm, reflecting difficulty in adapting to multiple contacts. Its lower disturbance was due to limited interaction, unlike RICE, which contacted and passed both branches. The

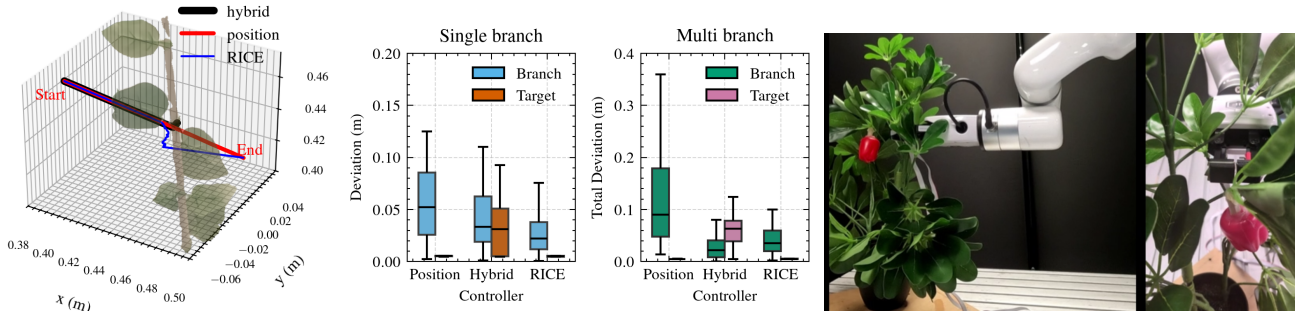

Fig. 6. Left: Example trajectories from three controllers interacting with a branch (experiment B). RICE (blue) moves around the branch to reach the target, the hybrid controller (black) stops on contact, and the position controller (red) pushes through. Centre: Branch deviation from OptiTrack and robot deviation from target for 20 single-branch trials and 10 multi-branch trials. Right: RICE controller in a highly cluttered environment, artificial plant setup, demonstrating its ability to maneuver around branches to reach the target.

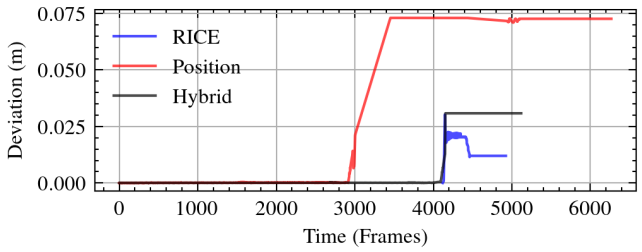

Fig. 7. Example of branch deviation recorded by the OptiTrack system. RICE (blue) causes minimal disturbance, the hybrid controller (black) pushes up to its 1 N threshold, and the position controller (red) generates the highest disturbance.

TABLE II

CONTROLLER PERFORMANCE FOR SINGLE AND MULTI-BRANCH SETUPS

Experiment	Controller	No breakage per trail	Target reached	No-Break Reach Rate
Single Branch	Hybrid	17/20 (broke 3 branches)	9/20	6/20 (30%)
	Position	15/20 (broke 5 branches)	**20/20**	15/20 (75%)
	RICE	**20/20** (broke 0 branches)	**20/20**	**20/20** (100%)
Multi Branch	Hybrid	**10/10** (broke 0 branches)	1/10	1/10 (10%)
	Position	1/10 (broke 12 branches)	**10/10**	1/10 (10%)
	RICE	**10/10** (broke 0 branches)	**10/10**	**10/10** (100%)

position-based controller followed a fixed path and broke branches in 9 out of 10 trials. It caused the most disturbance (median 90mm) but still reached the target in all cases. In Trial 1, it reached the target by bending the branches; in Trial 4, it broke both. This resulted in a 10% No-Break Reach rate, indicating poor adaptability in constrained spaces. These results highlight the advantages of RICE's responsive control strategies in cluttered plant environments, where adapting to contact and minimizing damage are critical.

D. RICE Reliability Evaluation

To further evaluate RICE, we ran 100 single-branch trials and 25 multi-branch trials. The position and hybrid controllers were excluded due to their consistent behavior, either reaching the target with disturbance or failing. RICE reached the target every time without breakage for both single branch setup (median deviation 22mm) and multi branch trials (median deviation 36mm; see Table III). Using a Bayesian estimate with a uniform prior [36], the expected probability of success in the next trial is

$$\mathbb{E}[\theta \mid R, N] = \frac{R+1}{N+2}, \tag{10}$$

commonly referred to as Laplace's rule of succession, where R is the number of successful trials out of N. This formulation provides a direct estimate of the probability of success in subsequent trials. Defining success as reaching the target without breakage, RICE achieved estimated probabilities of 0.99 in the single-branch setup and 0.96 in the multi-branch setup, demonstrating high reliability across the tested conditions.

TABLE III

PERFORMANCE OF RICE CONTROLLER FOR REPETITIVE TRIALS

Setup	No breakage per trial	Target Reached	No-Break Reach Rate
RICE (1 branch)	100/100	100/100	100/100
RICE (2 branches)	25/25 (broke 0 branch)	25/25	100/100

E. Artificial Plant

We evaluated all three controllers in a highly cluttered environment using an artificial plant setup. Each controller was tested with the same initial configuration in five different target locations. The RICE controller successfully moved around branches in all cases (Fig. 6 Right). The hybrid controller halted consistently upon contact, failing to progress toward the target. The position-based controller pushed through the branches, causing deformation or breakage. Representative behaviors and outcomes of these trials are shown in the supplementary media.

VI. DISCUSSION

Evaluating motion in dense, deformable environments like plant canopies is challenging. Real canopies are highly variable, with complex, unpredictable responses that are difficult to measure or replicate. To compare controllers fairly, we used a simplified setup that allowed consistent tracking of disturbance and performance. While it cannot fully capture real-world complexity, it provides a practical solution for systematic evaluation. We also tested on a realistic artificial plant to replicate clutter challenges. Together, these setups establish a baseline for future studies. These evaluations showed both RICE's adaptability and its limitations. RICE currently senses only at the tips of the gripper and operates on an industrial robot with limited dexterity for dense foliage. Algorithmically, it models the robot as a point mass, without

accounting for link collisions or environmental structures. An example of this was during one test a deformable leaf caused temporary repeated contact and motion. This edge case, influenced by trajectory direction and sensor location, resolved naturally as the plant deformed in different directions, producing varied outcomes. While RICE's tactile-only design provides local feedback for reactive motion, it lacks broader spatial awareness. Without a global view of plant structure, it cannot plan longer-term or more strategic pushing behaviours. Adding additional sensing, such as vision, could complement tactile input and support more informed, adaptive interactions in complex canopies.

VII. Conclusion

We introduced RICE, a reactive, model-free controller for safe motion in cluttered, contact-rich environments such as agricultural canopies. RICE combines real-time tactile feedback with end-effector position control in a hierarchical framework, enabling adaptive motion strategy through deformable foliage without an environmental model. This approach balances minimal disturbance by safely pushing through obstacles to reach occluded targets. In 30 trials with custom, trackable mock plants and five in denser foliage, RICE consistently reached targets with lower disturbance and no damage, clearly outperforming position and hybrid controllers in adaptability. These results demonstrate the potential of tactile-driven control for safe, effective autonomy in unstructured agricultural settings. Future work will extend RICE to outdoor environments and integrate additional sensing for improved performance.

References

[1] K. Legun and K. Burch, "Robot-ready: How apple producers are assembling in anticipation of new AI robotics," *Journal of Rural Studies*, 2021.

[2] M. R. Dogar and S. S. Srinivasa, "A planning framework for non-prehensile manipulation under clutter and uncertainty," *Autonomous Robots*, 2012.

[3] W. Liang, Q. Ren, X. Chen, J. Gao, and Y. Wu, "Dexterous manoeuvre through touch in a cluttered scene," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2021.

[4] A. Silwal, F. Yandun, A. K. Nellithimaru, T. Bates, and G. Kantor, "Bumblebee: A path towards fully autonomous robotic vine pruning." *Field Robotics*, vol. 2, 2022.

[5] T. Jin and X. Han, "Robotic arms in precision agriculture: A comprehensive review of the technologies, applications, challenges, and future prospects," *Computers and Electronics in Agriculture*, 2024.

[6] B. Siciliano and L. Villani, *Robot Force Control*. Springer Science & Business Media, 1999.

[7] M. Iskandar, C. Ott, A. Albu-Schäffer, B. Siciliano, and A. Dietrich, "Hybrid force-impedance control for fast end-effector motions," *IEEE Robot. Autom. Lett.*, vol. 8.

[8] S. Suresh, H. Qi, T. Wu, T. Fan, L. Pineda, M. Lambeta, J. Malik, M. Kalakrishnan, R. Calandra, M. Kaess *et al.*, "Neuralfeels with neural fields: Visuotactile perception for in-hand manipulation," *Science Robotics*, vol. 9, 2024.

[9] B. Frank, C. Stachniss, R. Schmedding, M. Teschner, and W. Burgard, "Real-world robot navigation amongst deformable obstacles," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2009.

[10] ——, "Learning object deformation models for robot motion planning," *Robotics and Autonomous Systems*, vol. 62, 2014.

[11] F. Grothe, V. N. Hartmann, A. Orthey, and M. Toussaint, "ST-RRT∗: Asymptotically-Optimal Bidirectional Motion Planning through Space-Time," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022.

[12] S. Rodríguez, J. M. Lien, and N. M. Amato, "Planning motion in completely deformable environments," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2006.

[13] R. Gieselmann and F. T. Pokorny, "Planning-augmented hierarchical reinforcement learning," *IEEE Robot. Autom. Lett.*, 2021.

[14] M. Missura, A. Roychoudhury, and M. Bennewitz, "Fast-replanning motion control for non-holonomic vehicles with aborting a," *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2022.

[15] J. Haviland and P. Corke, "NEO: A Novel Expeditious Optimisation Algorithm for Reactive Motion Control of Manipulators," *IEEE Robot. Autom. Lett.*, 2021.

[16] I. Rhee, G. Kang, S. J. Moon, Y. S. Choi, and H. R. Choi, "Hybrid impedance and admittance control of robot manipulator with unknown environment," *Intelligent Service Robotics*, 2023.

[17] M. Iskandar, A. Albu-Schäffer, and A. Dietrich, "Intrinsic sense of touch for intuitive physical human-robot interaction," *Science Robotics*, vol. 9, 2024.

[18] Z. Wang, L. Zou, X. Su, G. Luo, R. Li, and Y. Huang, "Hybrid force/position control in workspace of robotic manipulator in uncertain environments based on adaptive fuzzy control," *Robotics and Autonomous Systems*, vol. 145, 2021.

[19] T. Gold, A. Volz, and K. Graichen, "Model Predictive Interaction Control for Robotic Manipulation Tasks," *IEEE Transactions on Robotics*, 2023.

[20] M. D. Killpack, A. Kapusta, and C. C. Kemp, "Model predictive control for fast reaching in clutter," *Autonomous Robots*, 2016.

[21] Y. Xiong, Y. Ge, and P. J. From, "An obstacle separation method for robotic picking of fruits in clusters," *Computers and Electronics in Agriculture*, vol. 175, 2020.

[22] C. Lehnert, D. Tsai, A. Eriksson, and C. McCool, "3d move to see: Multi-perspective visual servoing towards the next best view within unstructured and occluded environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2019.

[23] J. Lee, C. Nam, J. Park, and C. Kim, "Tree Search-based Task and Motion Planning with Prehensile and Non-prehensile Manipulation for Obstacle Rearrangement in Clutter," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2021.

[24] S. Yao, S. Pan, M. Bennewitz, and K. Hauser, "Safe leaf manipulation for accurate shape and pose estimation of occluded fruits," *arXiv preprint arXiv:2409.17389*, 2024.

[25] E. Gursoy, D. Kulić, and A. Cherubini, "Occlusion handling by pushing for enhanced fruit detection," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 76–81.

[26] Z. Erickson, V. Gangaram, A. Kapusta, C. K. Liu, and C. C. Kemp, "Assistive Gym: A Physics Simulation Framework for Assistive Robotics," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 5 2020.

[27] J. Matas, S. James, and A. J. Davison, "Sim-to-Real Reinforcement Learning for Deformable Object Manipulation," *Conference on Robot Learning*, 2018.

[28] E. Tagliabue, A. Pore, D. Dall'Alba, E. Magnabosco, M. Piccinelli, and P. Fiorini, "Soft tissue simulation environment to learn manipulation tasks in autonomous robotic surgery," *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2020.

[29] J. Deng, S. Marri, J. Klein, W. Pałubicki, S. Pirk, G. Chowdhary, and D. L. Michels, "Gazebo plants: Simulating plant-robot interaction with cosserat rods," *arXiv preprint arXiv:2402.02570*, 2024.

[30] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The international journal of robotics research*, vol. 5, 1986.

[31] D. E. Whitney, "Resolved motion rate control of manipulators and human prostheses," *IEEE Transactions on man-machine systems*, vol. 10.

[32] UFactory, *xArm User Manual*, UFactory Robotics, 2024. [Online]. Available: https://www.ufactory.cc/lite-6-collaborative-robot/

[33] Xela Robotics. (2024) uskin patch sensors. [Online]. Available: https://www.xelarobotics.com/uskin-patch

[34] OptiTrack. (2025) Optitrack motion capture systems. [Online]. Available: https://www.optitrack.com/cameras/

[35] A. You, H. Kolano, N. Parayil, C. Grimm, and J. R. Davidson, "Precision fruit tree pruning using a learned hybrid vision/interaction controller," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022.

[36] C. Lehnert, C. McCool, I. Sa, and T. Perez, "Performance improvements of a sweet pepper harvesting robot in protected cropping environments," *Journal of Field Robotics*, vol. 37, no. 7, pp. 1197–1223, 2020.