

# Gazebo Plants: Simulating Plant-Robot Interaction with Cosserat Rods

Junchen Deng  
KAUST

junchen.deng@kaust.edu.sa

Wojtek Pałubicki  
AMU

wp06@amu.edu.pl

Samhita Marri  
UIUC

marri2@illinois.edu

Sören Pirk  
CAU

sp@informatik.uni-kiel.de

Jonathan Klein  
KAUST

jonathan.klein@kaust.edu.sa

Girish Chowdhary  
UIUC

girishc@illinois.edu

Dominik L. Michels  
KAUST

dominik.michels@kaust.edu.sa

## Abstract

*Robotic harvesting has the potential to positively impact agricultural productivity, reduce costs, improve food quality, enhance sustainability, and to address labor shortage. In the rapidly advancing field of agricultural robotics, the necessity of training robots in a virtual environment has become essential. Generating training data to automatize the underlying computer vision tasks such as image segmentation, object detection and classification, also heavily relies on such virtual environments as synthetic data is often required to overcome the shortage and lack of variety of real data sets. However, physics engines commonly employed within the robotics community, such as ODE, Simbody, Bullet, and DART, primarily support motion and collision interaction of rigid bodies. This inherent limitation hinders experimentation and progress in handling non-rigid objects such as plants and crops. In this contribution, we present a plugin for the Gazebo simulation platform based on Cosserat rods to model plant motion. It enables the simulation of plants and their interaction with the environment. We demonstrate that, using our plugin, users can conduct harvesting simulations in Gazebo by simulating a robotic arm picking fruits and achieve results comparable to real-world experiments.*

**Supplemental Material:** An accompanying video can be found on <https://youtu.be/r8Q31w4bNbs> showing the experiments presented in this paper.

**Plugin and Source Code:** Upon request, we are happy to share our GAZEBOPLANTS plugin open-source (MPL 2.0).

**Keywords:** Agricultural Robotics, Cosserat Rods, Gazebo, Plant Dynamics, Simulation, Virtual Environments.

## 1. Introduction

Robotics has demonstrated its efficacy across a broad spectrum of applications, ranging from the assembly of automobiles in factories using manipulator arms like those manufactured by KUKA, to streamlining household chores such as vacuuming through the utilization of mobile robots like iRobot’s Roomba. The structured nature of these artificially created environments has propelled significant advancements in research within these domains over the past few decades. However, the progress of robotics in agricultural applications has been comparatively less conspicuous, primarily due to the inherent unstructured and cluttered characteristics of agricultural environments. Nevertheless, there exists a pressing necessity for the advancement of agricultural robots to ensure food security for a growing global population in a sustainable manner [1], as well as to tackle issues like labor shortages.

While it is feasible to benchmark and assess different methods for a specific task within structured environments, such as picking a mug off a table, achieving a fair evaluation in agricultural settings is more challenging. Using the harvesting scenario as an illustration, several complexities arise. Firstly, no two plants, even of the same breed, possess identical structures. Secondly, once a fruit is harvested, reassessment becomes impractical. Thirdly, accessibility issues, such as cost or geography, can vary significantly among different plant varieties.

We believe that simulating plants will play an important role in addressing this disparity, akin to how progress in applications within structured environments was complemented by the development and accessibility of various simulation environments, as exemplified by [2, 3]. Simulation tools for robotics contribute significant value, not only

in the realm of robot learning but also in the crucial phase of verification before the deployment on physical robotic systems. This verification process is vital to ensuring the safety and functionality of both the robot and its operational environment. Particularly in data-intensive methods such as reinforcement learning, direct data collection on the robot may prove inefficient and potentially unsafe. Moreover, it encounters limitations in scalability, given the costs associated with employing multiple physical robots simultaneously to expedite the learning process. In contrast, simulation, as demonstrated in recent work like Issac Gym [4], allows several robots to learn in parallel and from each other, minimizing wear and tear on the robot and reducing risks to the environment. These simulation tools provide realistic sensing and testing capabilities that are not yet fully realized in agricultural settings.

The underlying computer vision tasks of agricultural robotics comprise, among others, image segmentation, object detection and classification. Addressing these tasks using state-of-the-art machine learning relies on the availability of sufficient training data. Such data can be generated in virtual environments resulting in large sets of synthetic data [5]. The use of synthetic training data in agricultural robotics has already been successfully demonstrated, for example, for the particular application of harvesting sweet pepper [6].

From a more general perspective, according to Gartner<sup>1</sup>, the majority of training data will anyhow be synthetic by the end of 2024. In particular, for automatizing computer vision tasks, the quality of these synthetic data sets is crucial and requires accurate simulations in virtual environments.

In this paper, we introduce a novel simulation framework for plants integrated as a plugin within the Gazebo system [7]. Our simulation framework not only encompasses simulated plants that dynamically respond to robot interactions but also incorporates the crucial functionality of the robot manipulator in detaching fruits. Furthermore, we conduct comprehensive calibration tests across various parameters to ensure a realistic and accurate simulation environment.

## 2. Related Work

The relevant prior work includes contributions to plant simulation, virtual environments for robots, and the application of robots in agriculture.

**Plant Simulation.** Numerical methods for the simulation of plant dynamics can be divided into two categories based on the discretization approach. The first one involves discretizing the plant into a mesh and using, for example, the Finite Element Method (FEM) [8, 9], while the second



Figure 1. Visual comparison of simulation and real-world scenario: UFactory Lite 6 robotic arm picking cherry tomato using our custom gripper (left) and our corresponding simulation in the virtual environment (right).

one involves discretizing it into a skeleton and, among others, using Position-Based Dynamics (PBD) [10].

To implement constraints between particles, PBD is often combined with Cosserat rods. This classical model has been introduced first in visual computing by [11] to simulate more physically accurate bending and twisting effects. The classical model as well as several modern adaptations have been widely utilized for the simulation of fiber-like objects such as threads and hair [11–20].

In addition to the dynamics of the plant itself, some research focuses on other natural phenomena related to plants. This includes aspects such as plant growth [21], the effect of water in plants on plant morphology [22, 23], the formation process of complex root architectures [24], the combustion of plants [25], the interaction between climate and vegetation [26], the influence of the environment on plant shape [27–31], the simulation of the cambium of trees [32], and modeling and simulating plant flowers [33, 34]. A botanical material model is proposed in [35] aiming for generating authentic simulations rooted in the biomechanics of trees. The development of complex ecosystems is simulated in [36].

**Virtual Environments for Robots.** There are various simulation platforms such as Webots [37] and MuJoCo [38] that are widely used in robotics. We chose Gazebo [7] to integrate our plugin as it comes with easy integration with ROS, one of the standard frameworks for deploying robotic systems. With additional support in ROS for various sensors such as vision and LIDAR, research in complex environments is made feasible. There has also been enormous development to support different plugins in Gazebo such as articulated rigid objects interacting with fluid [39], simulating radio frequency identification for industrial applications [40], and transceiver for rescue operations [41], which demonstrates its broad range of applications.

**Robotic Applications in Agriculture.** Previous contributions tackle various agricultural tasks using different robots like harvesting [42], phenotyping [43], and navigating in crop rows [44] by utilized simulation tools. In [42], sweet pepper plants are modeled in a virtual robot experimentation platform (VREP) to develop vision-based control

<sup>1</sup><https://techmonitor.ai/technology/ai-and-automation/ai-synthetic-data-edge-computing-gartner>

methods in reaching a target pepper. However, the detachment of the pepper is not simulated to perform the harvesting task. In the case of [43], plant volume and height are estimated in ROS-Gazebo where the CAD models of plants are developed in SketchUp. Using such synthetic data is more feasible than manual collection in the field where human errors can affect the ground truth. Autonomous navigation in crop rows is also studied in [43, 44] but both works are limited by the static nature of the CAD plant models as the dynamic behavior of plants, when external objects interact, is not simulated, causing a sim-to-real gap.

Other works include deploying artificial plant billboards like in [45], which capture the visual appearance of the plant well, but fail to mimic its dynamic motion. Another recent work [46] developed a physical proxy for grasping apples but it does not account for interactions with other parts of the plant. A more recent work [47] utilized Unity’s framework for the dynamic behavior of plants allowing robot-plant interaction in studying autonomous reaching of pruning targets, but the detachment of plant parts is not considered. Our proposed work not only allows for robot-plant interactions but also simulates the detachment of fruits, further motivating researchers to develop autonomous harvesting techniques.

**Paper Outline.** The remaining sections of the paper are structured as follows: Section 3 covers the methodology behind our plugin as well as implementation details such as the interaction between the plugin and the simulation platform Gazebo. Section 4 comprises experiments, encompassing the impact of plant parameters on plant behavior, experiments related to fruit harvesting, and a comparison between simulation results and real-world outcomes. Concluding remarks are provided in Section 5.

### 3. Methodology

#### 3.1. Interaction Framework

We developed our plant simulation plugin by utilizing the Gazebo World plugin as its foundation. Within Gazebo’s world file, users are required to specify the robot’s name, which will interact with the plant. In the initialization phase, the plugin then retrieves the geometric information of the robot object from Gazebo and subsequently constructs data structures for collision detection. It proceeds to generate the plant model using a combination of cylinders and spheres. The whole pipeline is shown in Figure 2. Given that we employ a position-based simulation method, the plugin only needs to periodically request the specific robot’s pose information from Gazebo at each time step. Since our implementation takes the form of a Gazebo plugin, users maintain the versatility to utilize our plugin with any library that possesses the capability to interface with Gazebo, such as ROS and MoveIt.

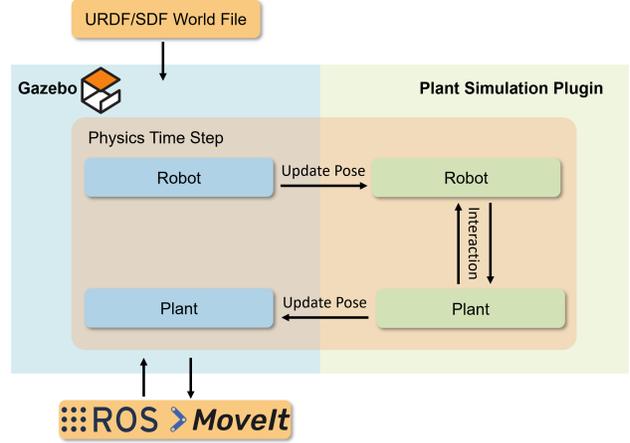


Figure 2. During initialization, Gazebo loads the world file, and within the parameters of our plugin, users need to specify both the plant file and the name of the robot intended to interact with the plant. Subsequently, the plugin proceeds to generate the plant as a combination of cylinders and spheres within the Gazebo environment, retrieving the robot’s geometric data from Gazebo. During each physics time step, the following sequence of actions unfolds: Firstly, the robot’s pose is updated, then it interaction with the plant is computed, and lastly, the pose of each individual component of the plant is reported back accordingly. Users can manipulate the robot by interfacing with Gazebo through widely used libraries such as ROS and MoveIt.

#### 3.2. Position and Orientation Based Cosserat Rods

We choose to adopt the Position-Based Dynamics (PBD) framework, as detailed in [48] to model plant branches as Cosserat rods. The Cosserat theory models a rod as the curve  $\mathbf{r}(s) : \mathbb{R} \rightarrow \mathbb{R}^3$  along the centerline. As shown on the left in Figure 3, for each point on the curve, there is an orthonormal frame with basis  $\{\mathbf{d}_1(s), \mathbf{d}_2(s), \mathbf{d}_3(s)\}$ , where  $\mathbf{d}_3(s)$  is the tangent direction of the curve. To establish this coordinate system, we rely on a quaternion  $q(s)$ , which serves as the means to rotate the basis vectors from the world coordinate system. The strain measure for bending and twisting is determined using the Darboux vector  $\Omega(s)$  which is defined as

$$\Omega(s) = 2\bar{q}(s) \frac{\partial}{\partial s} q(s),$$

where  $\bar{q}$  denotes the conjugate quaternion. Because some rods have an initial bending and torsion, we denote the rest pose Darboux vector as  $\Omega(s)^0$ . Thus, the strain measure of bending and torsion is defined as

$$\Delta\Omega(s) = \Omega(s) - \Omega(s)^0.$$

To discretize the curves, each rod element comprises two adjacent particles and a quaternion that holds both the position and orientation information of the coordinate system,

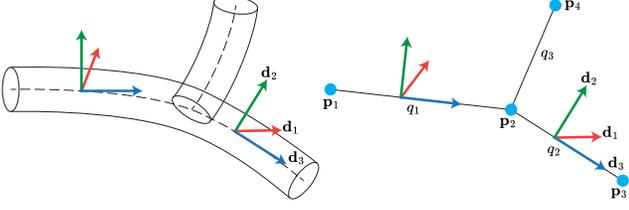


Figure 3. Illustration of the plan discretization in which  $\{\mathbf{d}_k\}$  denotes the local coordinate system for the Cosserat rod.

as shown in Figure 3. The Darboux vector should be also discretized between each neighbor rod. The derivative of the quaternion is discretized by a finite difference approximation  $\partial q_i / \partial s = (q_{i+1} - q_i) / l_i$ , where  $l_i$  denotes the average length of neighbor rods  $i$  and  $i + 1$ . Thus, the discrete Darboux vector is defined as

$$\Omega_i = \frac{2}{l_i} (\bar{q}_i q_{i+1}).$$

Here, to handle the ambiguity introduced by dropping the real part of the quaternion in [48], we employ the augmented Darboux vector [49] as above.

Within the PBD framework, we adopt a pair of constraints to enforce the desired behavior of the plant. Specifically, we make use of a stretch and shear constraint  $\mathbf{C}_s$ . This constraint is designed to approximately preserve the rod's length and relies on the configuration of two endpoints  $\mathbf{p}_1$  and  $\mathbf{p}_2$ , and on its length and original orientation  $\mathbf{d}_3$ :

$$\mathbf{C}_s(\mathbf{p}_1, \mathbf{p}_2, \mathbf{d}_3) = \frac{1}{l_0} (\mathbf{p}_2 - \mathbf{p}_1) - \mathbf{d}_3,$$

where  $l_0$  represents the rest length of each rod. To simulate bending and twisting behavior, we use constraint  $\mathbf{C}_b$ , which is defined based on the difference between the Darboux vector  $\Omega$  at the current pose and the rest pose:

$$\mathbf{C}_b(q_1, q_2) = \Omega - s\Omega_0,$$

$$s = \begin{cases} +1 & \text{for } |\Omega - \Omega_0|^2 < |\Omega + \Omega_0|^2, \\ -1 & \text{for } |\Omega - \Omega_0|^2 > |\Omega + \Omega_0|^2. \end{cases}$$

This adjustment of  $s$  is necessary to move the rod to the nearest rest pose since  $\Omega$  and  $-\Omega$  represent the same rotation.

### 3.3. Collision Detection and Response

In our simulation, we represent the rod as a capsule body, with radius information stored at each node. To simplify the modeling of fruits, we model them as spheres that are attached to a specific node.

Resolving collisions involves separating the two colliding objects by pushing them away from each other along the

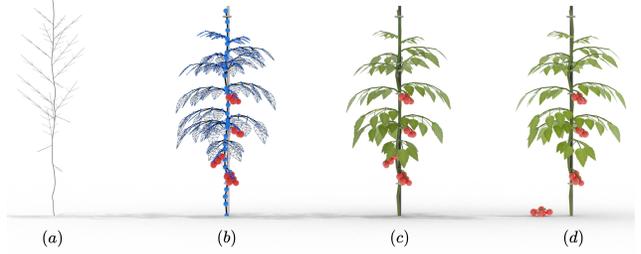


Figure 4. (a) portrays a plant through the use of multiple curves. Through the process of sampling along these curves, connecting the branches, and introducing leaves, we can transform it into (b), which represents a plant through a collection of particles linked by a graph. (c) depicts the plant, as represented by (b), in the physical world. (d) illustrates the action of harvesting a cluster of fruits from the plant.

normal vector at the point of collision. When dealing with self-collisions, the collision pairs include capsule-capsule, capsule-sphere, and sphere-sphere interactions. We calculate the collision position and its corresponding normal following the methodology outlined in [50].

We implement the robot/plant interaction as a one-way coupling, where the robot can move the plant but the plant does not affect the robots movement. This is a reasonable assumption for crop plants but may be less suitable for timber scenarios. The collision pairs in the crop context then include interactions between the sphere and the robotic arm, as well as between the capsule and the robotic arm. To handle collisions with the robotic arm, we utilize a Signed Distance Field (SDF)[51]. This approach enables us to efficiently determine the closest distance from any point to the robotic arm and obtain the corresponding normal direction. In the case of collisions involving the capsule body and the robotic arm, we performed equidistant sampling along the main axis of the capsule body. Subsequently, we identified the location of the collision as the point with the closest distance among all the sampled points.

### 3.4. Plant Fracture

Our framework supports plant fracture (Figure 4.d) through the assessment of strain. When a stretch or bending constraint value  $\mathbf{C}_s$  or  $\mathbf{C}_b$  exceeds a predefined threshold  $\mathbf{C}_s^{\max}$  or  $\mathbf{C}_b^{\max}$ , the corresponding rod element between the particles is fractured and removed from the system. Furthermore, any constraints associated with this particular rod are also eliminated.

### 3.5. Plant Model

We model plants by hand drawing curves for individual branches in 3D space (Figure 4.a). We then densely sample these curves using a step size smaller than the closest distance between any two curves. We establish a threshold,



Figure 5. Plant behaviors with different  $\sigma_{\text{distance}} = 0.02, 0.035, 0.05$  m. For all three results,  $\sigma_{\text{stiffness}} = 2 \times 10^4$  and  $\rho = 300 \text{ kg/m}^3$ .



Figure 6. Plant behaviors with different  $\sigma_{\text{stiffness}} = 1, 2, 3 \times 10^7$  Pa. For all three results,  $\sigma_{\text{distance}} = 0.02$  m and  $\rho = 300 \text{ kg/m}^3$ .

denoted as  $\sigma_{\text{connect}}$ , to add the branch connections. Nodes on one branch are connected to the endpoints of the nearest other branches when their distance is less than this threshold. This connection process is performed through breadth-first search, commencing from the specified root node, ensuring that the resulting connected graph forms a tree structure.

Next, we simplify the nodes within this tree by setting a distance threshold, denoted as  $\sigma_{\text{distance}}$ . Initially, certain nodes must be preserved, including junction nodes (nodes with a degree greater than 2) and leaf nodes (nodes with a degree equal to 1), referred to as key nodes. Subsequently, a depth-first search begins from the root node, and if the distance  $d$  from a node  $u$  to a preserved node  $v$  in its ancestry is less than  $\sigma_{\text{distance}}$ , then node  $u$  is also preserved. For key nodes, if their  $d$  is less than  $0.5\sigma_{\text{distance}}$  and  $v$  is not a key node, then node  $v$  is no longer retained. This approach prevents the issue of two retained nodes being too close to each other. Finally, we specify certain nodes for adding leaves and fruits, and we have predefined leaf templates for adding leaves.

Now, we represent the plant as a graph structure, where



Figure 7. Plant behaviors with different  $\rho = 300, 600, 900 \text{ kg/m}^3$ . For all three results,  $\sigma_{\text{distance}} = 0.02$  m and  $\sigma_{\text{stiffness}} = 2 \times 10^7$  Pa.

each edge in this graph corresponds to a rod, as illustrated in Figure 4.b. Even when dealing with leaves, rather than for a representation using individual triangles, we also treat their mesh as a graph. Each rod is subject to both stretch and shear constraints, and for every pair of adjacent rods, we add bend and twist constraints.

## 4. Experimental Results

### 4.1. Plant Parameters

We conduct several experiments to demonstrate that various plant parameters can influence the stiffness exhibited by plants. These parameters include  $\sigma_{\text{distance}}$ , used during the simplification of the plant model, the plant's stiffness parameter  $\sigma_{\text{stiffness}}$ , and the density of the plant's stem  $\rho$ .

In the Cosserat rod model, the stiffness of a plant is closely related to the number of nodes it possesses. When trees have a greater number of nodes, the system's degrees of freedom increase, resulting in an overall softer behavior of the plant. We have the ability to manage the number of nodes by modifying the distance parameter  $\sigma_{\text{stiffness}}$  during the tree creation process. As depicted in Figure 5, a smaller distance yields a greater number of nodes, causing the branches to exhibit more sag due to the influence of gravity.

In addition to the indirect influence of the number of nodes, we also have a direct parameter  $\sigma_{\text{stiffness}}$  that allows us to set Young's modulus at each node of the plant. This direct parameter gives us control over the stiffness of the plant. For herbaceous plants,  $\sigma_{\text{stiffness}}$  can be set in the order of  $10^6 \sim 10^7$  Pa, while for woody plants, it can be set in the order of  $10^8 \sim 10^{10}$  Pa. As depicted in Figure 6, from left to right, an increase in  $\sigma_{\text{stiffness}}$  results in the branches exhibiting a straighter behavior.

To align the plant with the target, the user will also need to specify the density of the plant's material. In the case of herbaceous plants, the typical density falls within the range

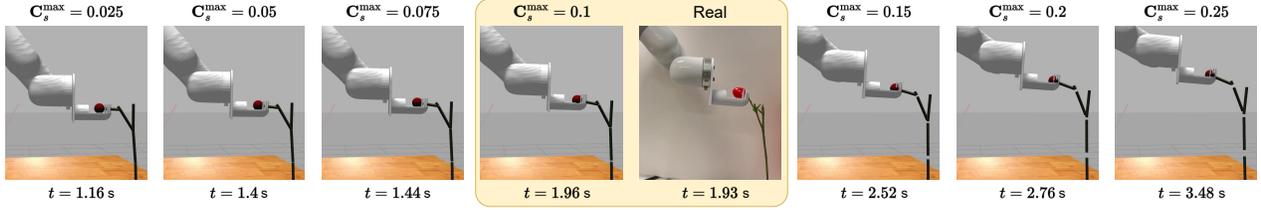


Figure 8. Performing a calibration experiment to identify the threshold for fruit detaching in Gazebo by comparing to a real demonstration where the robot moves upwards to detach the fruit. It is ensured that the manipulator’s trajectories in the real demonstration and simulation are similar. The snapshots shown here are taken just before the moment of detachment for various stress thresholds  $C_s^{\max}$  (labeled at the top) with the corresponding time taken to detach (labeled at the bottom). We observe that the detachment time in simulation is closest to the real demonstration at  $C_s^{\max} = 0.1$ , and as  $C_s^{\max}$  increases or decreases, the detachment time also increases or decreases respectively. Qualitatively, we also notice that the shape of the plant in simulation matches closest to the real demonstration at  $C_s^{\max} = 0.1$ . It is particularly noticeable at a higher  $C_s^{\max} = 0.25$ , where the plant in simulation stretches further than at  $C_s^{\max} = 0.1$ .

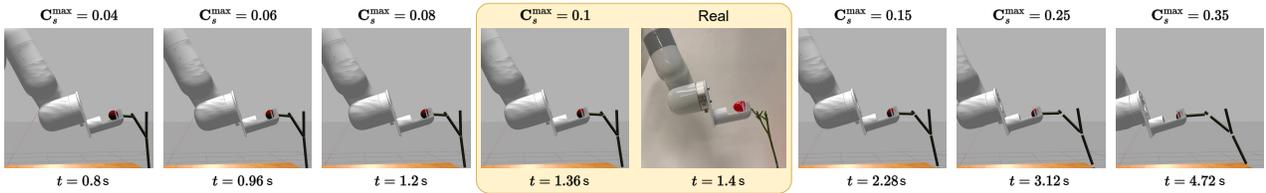


Figure 9. Performing a calibration experiment to identify the threshold for fruit detaching in Gazebo by comparing to a real demonstration where the robot moves horizontally away from the plant. It is ensured that the manipulator’s trajectories in the real demonstration and simulation are similar. The snapshots shown here are taken just before the moment of detachment for various stress thresholds  $C_s^{\max}$  (labeled at the top) with the corresponding time taken to detach (labeled at the bottom). We observe that the detachment time in simulation is closest to the real demonstration at  $C_s^{\max} = 0.1$ , and as  $C_s^{\max}$  increases or decreases, the detachment time also increases or decreases respectively. Qualitatively, we also notice that the shape of the plant in simulation matches closest to the real demonstration at  $C_s^{\max} = 0.1$ . It is particularly noticeable at a higher  $C_s^{\max} = 0.35$ , where the plant in simulation stretches further than at  $C_s^{\max} = 0.1$ .

of 200 to 1000 kg/m<sup>3</sup>. In Figure 7, we present the results with three distinct density values. A larger plant density leads to an increased bending under the influence of gravity and to a reduced susceptibility to external forces such as wind.

All experiments are conducted on a workstation with NVIDIA RTX 3090 GPU and Intel Xeon Gold 6136 3.00 GHz CPU. For  $\sigma_{\text{distance}} = 0.02$  m, the maximum number of rods in the scene is 2 242, with a 0.93 simulation to real-time ratio.

## 4.2. Calibration for Fruit Detachment

Since fruit detachment depends on the value of the strain constraint threshold  $C_s^{\max}$ , we perform a calibration using real world measurements. In particular, we perform two calibration experiments: The first one involves detaching the fruit in an upward motion of the robot (Figure 8), and the second one involves a horizontal motion away from the plant (Figure 9). For the demonstration of fruit detaching using the UFactory Lite 6 robotic arm, we used an artificial cherry tomato branch and a custom gripper as shown in Figure 1. We chose this gripper design over traditional scissors-based such as [52] to avoid causing damage to the

plant. We measure the time to detach the fruit from the start of the robot’s motion as reported at the bottom of each snapshot in Figure 8 and Figure 9.

In both cases, we notice that at  $C_s^{\max} = 0.1$ , the time taken to detach on the real system and in the simulation are similar. As seen in the figures, as the value of the stress threshold increases,  $C_s^{\max} > 0.1$ , there is an increase in the stretch of the plant, and for lower thresholds, the stretch is not as pronounced.

## 4.3. Comparison with Reality

To demonstrate the practicality of our simulator, we conducted a comparison between real-world outcomes and simulated results. Employing a robotic arm, we harvested a tomato in a greenhouse. Subsequently, we constructed a model of a similar plant and, by adjusting its parameters, successfully aligned the simulation results with the real-world outcomes in Figure 10.

As the tomato plant which we have harvested has been fixed from the upper end and suspended from a beam, attempting to pick the fruit by pulling would result in a significant displacement of the entire plant, exceeding the control range of the robotic arm. Consequently, we opted for a



Figure 10. Selected frames extracted from the real-world harvesting sequence (top row) demonstrate the actual tomato retrieval process using a robotic arm in a greenhouse. The corresponding frames at the bottom row depict our simulation results, showing the similarity between simulated and real-world harvesting scenarios.

bending method to harvest the fruit. During the simulation, we set the bending threshold between the two segments connected to the fruit  $C_b^{\max}$  to 0.31. The full experiment is shown in the accompanying video.

## 5. Conclusion

In this project, we have developed a plant simulator plugin for Gazebo, based on Position-Based Dynamics (PBD) and Cosserat rods. The plugin facilitates interaction between robots and plants, especially in the context of fruit harvesting. We support two methods for fruit picking: Stretching and bending. Compared to existing harvesting robots that use scissors, the stretching method closely resembles human picking, causing less damage to plants. On the other hand, the bending method allows for harvesting within smaller spaces. The comparison of simulation results with real-world outcomes validates the accuracy of our approach. We believe that our plugin enables users to train their harvesting robots, not limited to cherry tomatoes, for autonomous picking. This has the potential to enhance productivity and reduce labor costs by allowing robots to autonomously harvest a variety of crops.

While our plugin is a step towards realistic simulation of plants, there are limitations: The plant sags due to gravity instead of staying in the position we modeled. To stay in the initial position requires a large hardness value, which affects the dynamics of the plant; a better approach would be to model the plant to fall due to gravity, but this effect is not linear and requires the user to change the initial position of the plant several times. To solve this problem, the initial Darboux vector should be tuned by optimization.

## Acknowledgements

This work has been partially supported by KAUST through the baseline funding of the Computational Sciences Group.

## References

- [1] J. A. Foley, N. Ramankutty, K. A. Brauman, E. S. Cassidy, J. S. Gerber, M. Johnston, N. D. Mueller, C. O’Connell, D. K. Ray, P. C. West *et al.*, “Solutions for a cultivated planet,” *Nature*, vol. 478, no. 7369, pp. 337–342, 2011. 1
- [2] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016. 1
- [3] F. Xia, A. R. Zamir, Z.-Y. He, A. Sax, J. Malik, and S. Savarese, “Gibson env: real-world perception for embodied agents,” in *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE, 2018. 1
- [4] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac gym: High performance gpu-based physics simulation for robot learning,” 2021. 2
- [5] J. Klein, R. E. Waller, S. Pirk, W. Pałubicki, M. Tester, and D. Michels, “Synthetic data at scale: A paradigm to efficiently leverage machine learning in agriculture,” *Available at SSRN 4314564*, 2023. 2
- [6] C. W. Bac, J. Hemming, B. van Tuijl, R. Barth, E. Wais, and E. J. van Henten, “Performance evaluation of a harvesting robot for sweet pepper,” *Journal of Field Robotics*, vol. 34, no. 6, pp. 1123–1139, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21709> 2
- [7] N. Koenig and A. Howard, “Design and use paradigms for gazebo, an open-source multi-robot simulator,” in *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 3. IEEE, 2004, pp. 2149–2154. 2
- [8] Y. Zhao and J. Barbič, “Interactive authoring of simulation-ready plants,” *ACM Trans. Graph.*,

- vol. 32, no. 4, jul 2013. [Online]. Available: <https://doi.org/10.1145/2461912.2461961> 2
- [9] H. Wang, M. Kang, J. Hua, and X. Wang, “Modeling plant plasticity from a biophysical model: biomechanics,” 11 2013. 2
- [10] C. Deul, T. Kugelstadt, M. Weiler, and J. Bender, “Direct position-based solver for stiff rods,” *Computer Graphics Forum*, vol. 37, no. 6, pp. 313–324, 2018. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13326> 2
- [11] D. K. Pai, “Strands: Interactive simulation of thin solids using cosserat models,” *Computer Graphics Forum*, vol. 21, no. 3, pp. 347–352, 2002. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.00594> 2
- [12] J. Spillmann and M. Teschner, “CORDE: Cosserat Rod Elements for the Dynamic Simulation of One-Dimensional Elastic Objects,” in *Eurographics/SIGGRAPH Symposium on Computer Animation*, D. Metaxas and J. Popovic, Eds. The Eurographics Association, 2007.
- [13] D. L. Michels, J. P. T. Mueller, and G. A. Sobottka, “A physically based approach to the accurate simulation of stiff fibers and stiff fiber meshes,” *Computers & Graphics*, vol. 53, pp. 136–146, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0097849315001624>
- [14] D. L. Michels, V. T. Luan, and M. Tokman, “A stiffly accurate integrator for elastodynamic problems,” *ACM Trans. Graph.*, vol. 36, no. 4, jul 2017. [Online]. Available: <https://doi.org/10.1145/3072959.3073706>
- [15] F. Bertails, B. Audoly, M.-P. Cani, B. Querleux, F. Leroy, and J.-L. L  v  que, “Super-Helices for Predicting the Dynamics of Natural Hair,” *ACM Transactions on Graphics*, 2006. [Online]. Available: <https://inria.hal.science/inria-00384718>
- [16] D. A. Lyakhov, V. P. Gerdt, A. G. Weber, and D. L. Michels, “Symbolic-numeric integration of the dynamical cosserat equations,” in *Computer Algebra in Scientific Computing*, V. P. Gerdt, W. Koepf, W. M. Seiler, and E. V. Vorozhtsov, Eds. Cham: Springer International Publishing, 2017, pp. 301–312.
- [17] D. L. Michels, G. A. Sobottka, and A. G. Weber, “Exponential integrators for stiff elastodynamic problems,” *ACM Trans. Graph.*, vol. 33, no. 1, feb 2014. [Online]. Available: <https://doi.org/10.1145/2508462>
- [18] H. Shao, T. Kugelstadt, T. H  drich, W. Pa  bicki, J. Bender, S. Pirk, and D. L. Michels, “Accurately solving rod dynamics with graph learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [Online]. Available: <http://computationalsciences.org/publications/shao-2021-physical-systems-graph-learning.html>
- [19] H. Shao, “Fluids, threads and fibers: towards high performance physics-based modeling and simulation,” Ph.D. dissertation, 2022.
- [20] J. Spillmann and M. Teschner, “Cosserat nets,” *IEEE transactions on visualization and computer graphics*, vol. 15, no. 2, pp. 325–338, 2008. 2
- [21] S. Pirk, T. Niese, O. Deussen, and B. Neubert, “Capturing and animating the morphogenesis of polygonal tree models,” *ACM Trans. Graph.*, vol. 31, no. 6, nov 2012. [Online]. Available: <https://doi.org/10.1145/2366145.2366188> 2
- [22] S. Jeong, S.-H. Park, and C.-H. Kim, “Simulation of Morphology Changes in Drying Leaves,” *Computer Graphics Forum*, 2013. 2
- [23] F. Maggioli, J. Klein, T. Haedrich, E. Rodol  , W. Pa  bicki, S. Pirk, and D. L. Michels, “A physically-inspired approach to the simulation of plant wilting,” in *SIGGRAPH Asia 2023 Conference Papers*, ser. SA ’23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi.org/10.1145/3610548.3618218> 2
- [24] B. Li, J. Klein, D. L. Michels, B. Benes, S. Pirk, and W. Pa  bicki, “Rhizomorph: The coordinated function of shoots and roots,” *ACM Trans. Graph.*, vol. 42, no. 4, jul 2023. [Online]. Available: <https://doi.org/10.1145/3592145> 2
- [25] S. Pirk, M. Jarzabek, T. Haedrich, D. L. Michels, and W. Pa  bicki, “Interactive wood combustion for botanical tree models,” *ACM Trans. Graph.*, vol. 36, no. 6, nov 2017. [Online]. Available: <https://doi.org/10.1145/3130800.3130814> 2
- [26] W. Pa  bicki, M. Makowski, W. Gajda, T. Haedrich, D. L. Michels, and S. Pirk, “Ecoclimates: Climate-response modeling of vegetation,” *ACM Trans. Graph.*, vol. 41, no. 4, jul 2022. [Online]. Available: <https://doi.org/10.1145/3528223.3530146> 2
- [27] S. Pirk, O. Stava, J. Kratt, M. A. M. Said, B. Neubert, R. M  ch, B. Benes, and O. Deussen, “Plastic trees: Interactive self-adapting botanical tree models,” *ACM Trans. Graph.*, vol. 31, no. 4, jul 2012. [Online]. Available: <https://doi.org/10.1145/2185520.2185546> 2
- [28] R. Habel, A. Kusternig, and M. Wimmer, “Physically guided animation of trees,” *Computer Graphics Forum*, vol. 28, no. 2, pp. 523–532, 2009. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01391.x>
- [29] E. Quigley, Y. Yu, J. Huang, W. Lin, and R. Fedkiw, “Real-time interactive tree animation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 5, pp. 1717–1727, 2018.

- [30] S.-K. Wong and K.-C. Chen, “A procedural approach to modelling virtual climbing plants with tendrils,” *Computer Graphics Forum*, vol. 35, no. 8, pp. 5–18, 2016. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12736>
- [31] T. Hädrich, B. Benes, O. Deussen, and S. Pirk, “Interactive modeling and authoring of climbing plants,” *Computer Graphics Forum*, vol. 36, no. 2, pp. 49–61, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13106>
- [32] J. Kratt, M. Spicker, A. Guayaquil, M. Fiser, S. Pirk, O. Deussen, J. C. Hart, and B. Benes, “Woodification: User-Controlled Cambial Growth Modeling,” *Computer Graphics Forum*, 2015. 2
- [33] A. Owens, M. Cieslak, J. Hart, R. Classen-Bockhoff, and P. Prusinkiewicz, “Modeling dense inflorescences,” *ACM Trans. Graph.*, vol. 35, no. 4, jul 2016. [Online]. Available: <https://doi.org/10.1145/2897824.2925982>
- [34] L. Ringham, A. Owens, M. Cieslak, L. D. Harder, and P. Prusinkiewicz, “Modeling flower pigmentation patterns,” *ACM Trans. Graph.*, vol. 40, no. 6, dec 2021. [Online]. Available: <https://doi.org/10.1145/3478513.3480548>
- [35] B. Wang, Y. Zhao, and J. Barbič, “Botanical materials based on biomechanics,” *ACM Trans. Graph.*, vol. 36, no. 4, jul 2017. [Online]. Available: <https://doi.org/10.1145/3072959.3073655>
- [36] M. Makowski, T. Hädrich, J. Scheffczyk, D. L. Michels, S. Pirk, and W. Pałubicki, “Synthetic silviculture: multi-scale modeling of plant ecosystems,” *ACM Trans. Graph.*, vol. 38, no. 4, jul 2019. [Online]. Available: <https://doi.org/10.1145/3306346.3323039>
- [37] O. Michel, “Webotstm: Professional mobile robot simulation,” 2004. 2
- [38] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033. 2
- [39] E. Angelidis, J. Bender, J. Arreguit, L. Gleim, W. Wang, C. Axenie, A. Knoll, and A. Ijspeert, “Gazebo fluids: Sph-based simulation of fluid interaction with articulated rigid body dynamics,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 238–11 245. 2
- [40] A. A. Alajami, G. Moreno, and R. Pous, “A ros gazebo plugin design to simulate rfid systems,” *IEEE Access*, vol. 10, pp. 93 921–93 932, 2022. 2
- [41] J. Cacace, N. Mimmo, and L. Marconi, “A ros gazebo plugin to simulate arva sensors,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7233–7239. 2
- [42] R. R. Shamshiri, I. A. Hameed, M. Karkee, and C. Weltzien, “Robotic harvesting of fruiting vegetables: A simulation approach in v-rep, ros and matlab,” *Proceedings in automation in agriculture-securing food supplies for future generations*, vol. 126, pp. 81–105, 2018. 2
- [43] J. Iqbal, R. Xu, S. Sun, and C. Li, “Simulation of an autonomous mobile robot for lidar-based in-field phenotyping and navigation,” *Robotics*, vol. 9, no. 2, p. 46, 2020. 2, 3
- [44] G. B. P. Barbosa, E. C. Da Silva, and A. C. Leite, “Vision-based autonomous crop row navigation for wheeled mobile robots using super-twisting sliding mode control,” in *2021 European Conference on Mobile Robots (ECMR)*, 2021, pp. 1–6. 2, 3
- [45] F. Taqi, F. Al-Langawi, H. Abdulraheem, and M. El-Abd, “A cherry-tomato harvesting robot,” in *2017 18th International Conference on Advanced Robotics (ICAR)*, 2017, pp. 463–468. 3
- [46] A. Velasquez, N. Swenson, M. Cravetz, C. Grimm, and J. R. Davidson, “Predicting fruit-pick success using a grasp classifier trained on a physical proxy,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 9225–9231. 3
- [47] Q. T. Barthelme and C. Lehnert, “Robotic crop handling in cluttered and unstructured environments using simulated l-system dynamic plant models,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 10 088–10 095. 3
- [48] T. Kugelstadt and E. Schömer, “Position and Orientation Based Cosserat Rods,” in *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, L. Kavan and C. Wojtan, Eds. The Eurographics Association, 2016. 3, 4
- [49] J. Hsu, T. Wang, Z. Pan, X. Gao, C. Yuksel, and K. Wu, “Sag-free initialization for strand-based hybrid hair simulation,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2023)*, vol. 42, no. 4, 07 2023. [Online]. Available: <https://doi.org/10.1145/3592143>
- [50] C. Ericson, *Real-Time Collision Detection*. USA: CRC Press, Inc., 2004. 4
- [51] D. Koschier *et al.*, “Discregrid Library.” [Online]. Available: <https://github.com/InteractiveComputerGraphics/Discregrid>
- [52] J. Jun, J. Kim, J. Seol, J. Kim, and H. I. Son, “Towards an efficient tomato harvesting robot: 3d perception, manipulation, and end-effector,” *IEEE access*, vol. 9, pp. 17 631–17 640, 2021. 6