

# Occlusion Reasoning for Skeleton Extraction of Self-Occluded Tree Canopies

Chung Hee Kim and George Kantor

**Abstract**—In this work, we present a method to extract the skeleton of a self-occluded tree canopy by estimating the unobserved structures of the tree. A tree skeleton compactly describes the topological structure and contains useful information such as branch geometry, positions and hierarchy. This can be critical to planning contact interactions for agricultural manipulation, yet is difficult to gain due to occlusion by leaves, fruits and other branches. Our method uses an instance segmentation network to detect visible trunk, branches, and twigs. Then, based on the observed tree structures, we build a custom 3D likelihood map in the form of an occupancy grid to hypothesize on the presence of occluded skeletons through a series of minimum cost path searches. We show that our method outperforms baseline methods in highly occluded scenes, demonstrated through a set of experiments on a synthetic tree dataset. Qualitative results are also presented on a real tree dataset collected from the field.

## I. INTRODUCTION

With increasing global population and labor shortages, modern agriculture is adopting new technologies to enhance sustainability and profitability. There is growing effort in developing robotic solutions to automate repetitive and laborious tasks that often require complex and delicate interaction with crops. Harvesting and pruning, for example, may require the agent to manipulate crops by pushing aside leaves or branches to reach occluded regions before picking or cutting. Understanding the unstructured and cluttered task environment is critical to automating such challenging tasks.

In this work, we address the perception problem, where robots must be able to sense and understand the complex task environment prior to crop interaction. We are particularly interested in extracting the skeleton of a self-occluded tree canopy by estimating the unobserved parts of the tree. A tree skeleton is useful as it describes the topological structure and contains useful information such as branch dimensions, positions, and hierarchy. This knowledge is helpful in phenotyping crop characteristics for growth assessment and can be critical for planning contact interactions, such as determining optimal pruning locations or generating trajectories to pick fruits by pulling a branch into its workspace.

One of the main challenges faced during the extraction of the tree skeleton from sensor data (typically in the form of point clouds obtained from laser scanners or depth cameras) stems from noise and occlusion, which is exacerbated by the presence of foliage as depicted in Fig. 1(a). Although there exists prior works that address skeletonization of occluded tree canopies, they are tailored towards sparse point clouds

C. H. Kim and G. Kantor are with Carnegie Mellon University, Pittsburgh PA, USA {chunghek, kantor}@andrew.cmu.edu



Fig. 1. (a) The 3D reconstruction (middle) and skeletonization (right) of a heavily occluded tree canopy (left). (b) Tree crops are tightly organized in rows in industrial orchard settings.

obtained from terrestrial laser scanners (TLS) and can be ineffective when applied to dense point clouds acquired from RGB-D or stereo cameras. Furthermore, it is generally assumed in prior works that the tree point cloud to be skeletonized is pre-registered from a 360 degree scan, which is difficult to obtain in industrial orchard settings where trees are tightly organized in rows as shown in Fig. 1(b).

We present a novel skeletonization method that extracts tree skeletons from one-sided views of the tree canopy by using a depth camera attached to a robot arm. Our method particularly addresses the challenge of skeletonizing heavily occluded tree canopies by observing that branches in nature generally extend linearly. We use this as a heuristic assumption to approximate a 3D likelihood map in the form of an occupancy grid to predict presence of occluded branch structures by searching for minimum cost paths. We validate the effectiveness of our approach by presenting quantitative assessment on a synthetic tree dataset with known ground truth, and present qualitative results from a real tree dataset collected at an apple orchard.

## II. RELATED WORK

A skeleton can be defined as a curve expressing the shape of an object that is consistent with the topological structure as well as the connectivity of the original object shape. Many methods have been proposed to extract skeletons from unordered 3D point clouds. A skeleton can be extracted by measuring the L1-median to determine local centers of the point cloud [1]. A Laplacian-based contraction method is proposed in [2], which collapses the input point cloud into a minimal volume using an iterative Laplacian smoothing process. In [3], a skeleton is extracted by defining a new feature representation called rotational symmetry axis. The aforementioned methods work well with limited noise and occlusion, while we are particularly interested in skeletonizing point cloud of botanical trees which are often noisy and

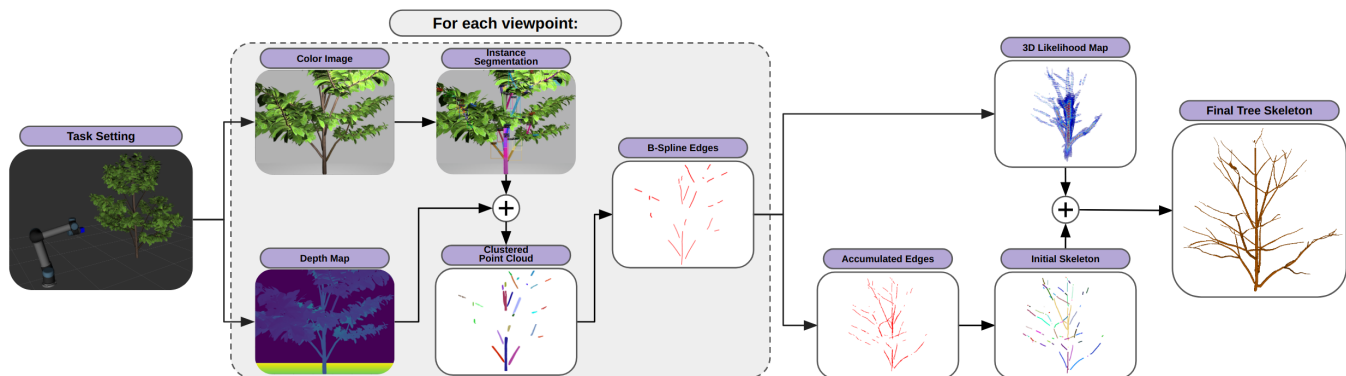


Fig. 2. Our system pipeline for tree skeletonization. It takes as input a series of RGB-D images of a self-occluded tree from multiple viewpoints and outputs the underlying tree skeleton.

occluded.

There is particular interest in extracting the skeleton of a tree canopy due to implications in automating agricultural tasks [4] such as pruning [5], harvesting [6], or phenotyping to estimate crop characteristics for growth assessment [7]. Prior knowledge about tree structures are commonly used to extract tree skeletons, such as branching properties [8], cylindrical shape priors [9] and upright offshoots [10]. Tree skeletons are commonly represented as graphs composed of vertices and edges; [11] and [12] uses graph-based methods to extract tree skeletons. A geometry-based method is used to fit cylinders on the point cloud in a hierarchical data structure to encapsulate parent-child relations of branches [13].

While the above mentioned tree skeletonization methods are tailored for trees without leaves, we address a more challenging problem in which tree canopies are self-occluded by leaves, fruits, or other branches. To address the problem of missing data, [14] proposes an iterative data-completion method to recover data for 3D tree modelling. Tree point clouds with leaves are often first separated into branches and leaves before skeleton extraction [15], while [16] directly generates visually convincing tree skeletons based on optimization driven by biological priors. Several works improve skeleton connectivity by joining disconnected skeletal structures through geometry-based methods [17], [18]; we use it as a baseline method and quantitatively show that our method outperforms [17] in occluded scenes. Despite addressing the issue of occlusion, these methods are built for sparse TLS point clouds and are inapplicable to dense point clouds. More importantly, it is assumed that the tree point cloud to be skeletonized is pre-registered from a 360 degree scan, which may be unpractical in industrial orchard setting where trees are organized in rows. The key contributions of this paper are:

- A novel skeletonization method that extracts tree skeletons from one-sided views of the tree canopy;
- A method to reason about the unobserved structures of the tree to predict the presence of occluded skeletons;
- Experimental results on a synthetic dataset quantitatively compared against existing baselines, as well as qualitative results on a real tree dataset collected from the field.

### III. METHODOLOGY

#### A. System Overview

Our tree skeletonization method aims to improve the knowledge of occluded regions in the tree that is often self-occluded by foliage, fruits, or other branches. Driven by a heuristic assumption on tree structures, we approximate a 3D likelihood map in the form of an occupancy grid that stores information of visible as well as occluded structures of a tree canopy to extract its skeleton. Fig. 2 shows an overview of our framework. The system takes as input a series of color and depth images from varying viewpoints obtained from a camera attached to a robot arm with known poses. The images are passed onto an instance segmentation network to generate a semantic point cloud of branch clusters (Sec. III-B). The occupancy probability distribution of the 3D likelihood map is updated based on the observed branch clusters over a sequence of images (Sec. III-C). The final skeleton is generated by joining disconnected branch skeletons through a series of minimum cost path searches in the likelihood map (Sec. III-D).

#### B. Semantic Point Cloud Acquisition

We first acquire a 3D point cloud semantically labeled with visible parts of the tree trunk, branches, and twigs (collectively referred to as branches from this point on) that are not occluded. This is achieved by projecting branch segmentation masks in 2D color images on to the 3D point cloud obtained from the depth image. The points corresponding to branches in the semantic point cloud are clustered based on the projected masks.

Our branch segmentation is based on the Mask R-CNN [19] instance segmentation network with a Feature Pyramid Network and ResNet50 backbone. The model takes as input  $1440 \times 1080$  images and outputs instance segmentation masks as well as its confidence scores ranging from 0 to 1. The network was trained on 130 manually labeled images (105 real images and 25 synthetic images) for 2000 iterations on an NVIDIA GeForce RTX 3070. Sample segmentation results are depicted in Fig. 3(a). In order to perform line-fitting on the point clusters (further described in Sec. III-C), we deliberately labeled each branch instances in the training images to be a slender polygon such that the projected




Fig. 3. (a) Instance segmentation results on a real tree image (left) and a synthetic tree image (right). (b) Tree forks (left) or angled branch instances (right) are labeled as multiple branch instances in the training dataset.

branch point cloud cluster in 3D space is also piece-wise slender. For example, a tree fork or a branch instance with a significant change in direction is labeled as two or more branch instances instead of a single instance as shown in Fig. 3(b).

### C. Skeleton Occupancy Likelihood Map

We now propose our novel approach to hypothesize on the presence of occluded branches by approximating a 3D likelihood map in the form of an occupancy grid, where the  $i$ -th grid voxel  $\mathbf{m}_i$  stores the skeleton occupancy probability  $p_\ell(\mathbf{m}_i) \in [0, 1]$ . The approximation is driven by a heuristic assumption: Given a visible branch instance, it is likely that the branch structure extends longer along its growth direction based on the observation that branches in nature generally grow straight. We model this as individually observed probability distributions  $p_o$  with an ellipsoidal contour obtained from the branch point cloud clusters as follows.

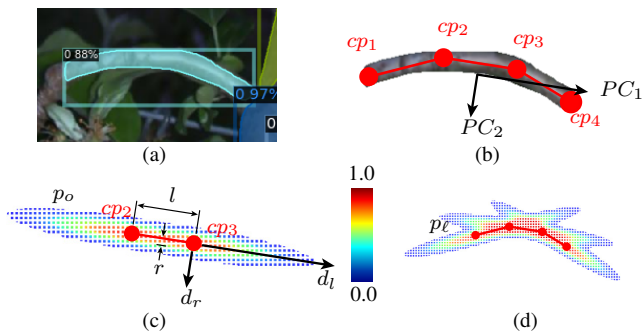


Fig. 4. (a) A single instance of a segmented branch component with a confidence score of 0.88. (b) The branch instance is projected into 3D space as a dense point cloud cluster. We compute the principal components ( $PC_i$ ) and fit a B-spline curve with four control points ( $cp_i$ ) on the point cloud cluster. (c)  $p_o$  is obtained from each line segment, (d) which updates  $p_\ell$  using equation (2).

A 3-dimensional B-Spline curve of degree-one with four control points is approximated on a point cloud cluster using the least squares method [20]. This results in three connected line segments that best represent the skeletal geometry of the point cloud cluster as shown in Fig. 4(b), which are




Fig. 5. Visualization of the overall joint 3D likelihood map (right) obtained from accumulated line segments (middle), which were extracted from observed branch clusters over the sequence of images (left).

accumulated over the sequence of images (Fig. 5). The radius of the branch instance is also estimated by performing principal component analysis on the point cloud cluster (see Fig. 4(b)), where the distance between the maxima and minima points along the second principal component is approximated to be the diameter of the branch instance. The probability distribution of the joint likelihood map  $p_\ell$  is updated by  $p_o$  based on the following rules:

- 1)  $p_o$  of the voxels containing the line segment is equal to the mask confidence score obtained from the instance segmentation network.
- 2)  $p_o$  of the voxels in the proximity of the line segment diminishes in the axial and radial directions along an ellipsoidal contour (Fig. 4(c)):

$$p_o = c - \frac{c}{k} \sqrt{\left(\frac{d_l}{l}\right)^2 + \left(\frac{d_r}{r}\right)^2} \quad (1)$$

where  $c$  is the mask confidence score,  $l$  is the line segment length, and  $r$  is the estimated radius.  $d_l$  and  $d_r$  is the axial and radial distance of the voxel from the line segment, and  $k$  is a parameter that controls the rate at which  $p_o$  decreases as  $d_l$  and  $d_r$  increases. For example, a large  $k$  results in a small decreasing rate, effectively enlarging the size of the ellipsoid. We empirically set  $k = 3$  in our experiments.

- 3) The voxel  $\mathbf{m}_i$  of the 3D likelihood map is updated by the observed occupancy probability  $p_o$  according to the following equation (Fig. 4(d)):

$$p_\ell(\mathbf{m}_i) = 1 - [1 - p_\ell(\mathbf{m}_i)][1 - p_o(\mathbf{m}_i)] \quad (2)$$

The resulting joint likelihood map updated by  $p_o$  from all line segments is depicted in Fig. 5. The 3D likelihood map effectively has higher occupancy probability in voxels that are directly intersected by a line segment, as well as in voxels that are jointly influenced by more than one line segment.

### D. Skeleton Extraction

The final tree skeleton  $G_T = (V_T, E_T)$  is an undirected acyclic graph with vertices  $V_T$  and edges  $E_T$  representing the topological structure of the tree canopy. This is obtained by post-processing the accumulated line segments with the joint likelihood map as follows.

We first obtain an initial tree skeleton  $G_{\text{init}} = (V_{\text{init}}, E_{\text{init}})$  by consolidating the accumulated line segments. To do so,

all line segments are converted into a set of vertices by sampling points equally spaced along the line (in practice, we set the spacing to be the voxel size of the likelihood map). The vertices are consolidated into a skeletal curve using the Laplacian smoothing method which repositions each vertex to the mean of its  $k$ -nearest neighbors:

$$v_i = \frac{1}{k} \sum_{j=1}^N v_j \quad (3)$$

where  $k$  is the number of neighborhood vertices,  $v_j$  is the position of neighbor vertex  $j$ , and  $v_i$  is the new position of vertex  $i$ . Laplacian smoothing is iteratively applied until the total change in vertex position per iteration converges below a threshold, which results in a set of well-refined skeletal vertices  $V_{\text{init}}$ . The edges  $E_{\text{init}}$  are initialized by building a Euclidean minimum spanning tree from  $V_{\text{init}}$  with edges that are longer than the voxel size removed.




Fig. 6. (a) The initial skeleton  $G_{\text{init}}$  shown as the curves colored black overlapped with the weighted likelihood graph  $G_\ell$ .  $G_{\text{init}}$  is composed of disjoint subgraphs  $g_i$  due to occlusion. (b) We search for minimum cost paths in  $G_\ell$  to join disconnected subgraphs  $g_i$ .

Due to occlusion by leaves and branches,  $G_{\text{init}} = \{g_i\}$  is composed of disjoint skeletal subgraphs  $g_i$  of a single tree. In order to predict the occluded parts of the tree, we join maximally connected subgraphs in  $G_{\text{init}}$  by searching for minimum cost paths in the 3D likelihood map as illustrated in Fig. 6 and summarized in Algorithm 1. The likelihood map is first converted into a weighted graph  $G_\ell$  by adding undirected edges between all  $3 \times 3 \times 3$  adjacent nodes (Line 2, Algorithm 1). The cost of an edge connecting vertex  $u$  and  $v$  is set to be the negative log of the average occupancy probability:

$$c(e_{uv}) = -\log\left(\frac{p_\ell(\mathbf{m}_u) + p_\ell(\mathbf{m}_v)}{2}\right) \quad (4)$$

which results in low cost for edges between nodes with high occupancy probability. Using the weighted likelihood graph  $G_\ell$ , the minimum cost path between all pairwise subgraph combinations in  $G_T$  (initialized with  $G_{\text{init}}$ ) are computed (Line 5-7, Algorithm 1). The path with the minimum cost is added to  $G_T$  (Line 9-10, Algorithm 1). This process is repeated until no more paths are found resulting in our final skeleton output  $G_T$ , composed of skeletal curves that were directly observed, as well as predicted skeletal curves that were unobservable due to occlusion.

---

#### Algorithm 1 Likelihood Map Path Search

---

**Input:** Initial Skeleton  $G_{\text{init}}$ , Likelihood Map  $p_\ell$   
**Output:** Tree Skeleton  $G_T$

- 1:  $G_T \leftarrow G_{\text{init}}$
- 2:  $G_\ell \leftarrow \text{WeightedUndirectedGraph}(p_\ell)$
- 3: **Repeat**
- 4:    $\mathcal{P} \leftarrow \text{EmptyArray}()$
- 5:   **for all** pairwise subgraph combination  $(g_u, g_v) \in G_T$  **do**
- 6:      $\rho \leftarrow \text{DijkstrasMinPath}(g_u, g_v, G_\ell)$
- 7:      $\mathcal{P}.\text{append}(\rho)$
- 8:   **end for**
- 9:    $\rho_{\text{min}} \leftarrow \text{MinCostPath}(\mathcal{P})$
- 10:  $G_T \leftarrow \text{JoinSubgraphs}(G_T, \rho_{\text{min}})$
- 11: **Until** no more paths are found

---

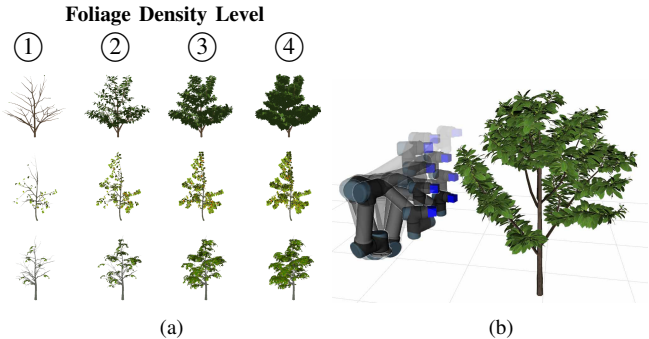


Fig. 7. (a) Synthetic meshes of an oak tree (top row), apple tree (middle row), and a walnut tree (bottom row) with varying foliage density increasing from left to right. Trees in the same row have the same structural topology. Our synthetic dataset consists of ten unique topology per tree species. (b) The robot is controlled to take images of the tree canopy from 10 different viewpoints in the Gazebo simulation environment.

## IV. EXPERIMENTS

We evaluate our proposed tree skeletonization method through quantitative evaluation and visual assessment. For the former, we collected a synthetic tree mesh dataset with known ground truth skeleton and propose metrics to measure the precision of the skeleton as well as the effectiveness of predicting unseen branch skeletons. Visual assessment is presented for the simulated dataset as well as a real tree dataset collected from an apple orchard.

### A. Experiments on Synthetic Trees

Experiments in simulation were performed on three different types of synthetic trees including oak, apple, and walnut trees. For each species, ten different structural topologies with four levels of varying foliage density were generated as depicted in Fig. 7(a), totalling to 120 trees. The synthetic tree meshes were created in SpeedTree<sup>1</sup> which were imported into a Gazebo<sup>2</sup> simulation environment to model an RGB-D camera attached to a UR5 robot arm. For each tree, the robot is controlled to collect images from 10 different viewpoints as shown in Fig. 7(b), which are sequentially passed onto our skeletonization pipeline.

We propose the following metrics to assess the correctness of the constructed skeleton as well as the effectiveness of our approach:

<sup>1</sup>www.speedtree.com

<sup>2</sup>www.gazebosim.org

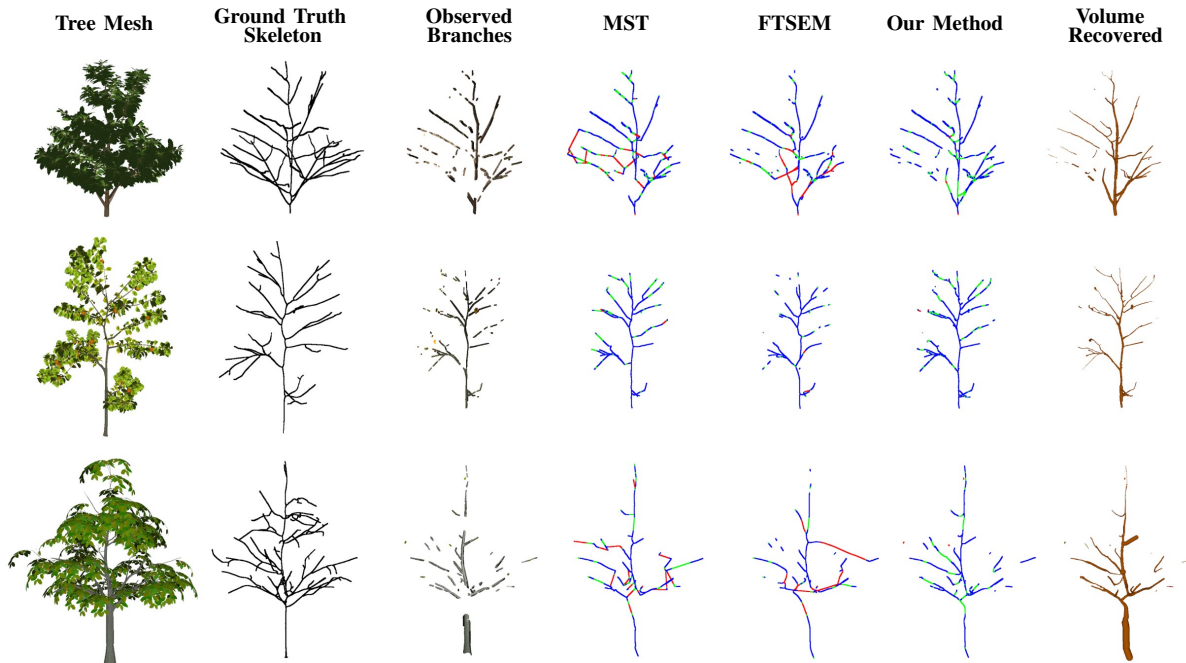


Fig. 8. Experiment results of synthetic trees including oak (top row), apple (middle row), and walnut (bottom row) trees with foliage density level ③. We compare our method against MST-based and FTSEM-based methods. The extracted skeleton is color coded:  $TP$ ,  $TP_{occ}$ , and  $FP$  vertices are colored blue, green, and red, respectively. The right-most column depicts the skeleton extracted from our method with recovered volume based on the radius information of the skeletal vertices.

- **Skeleton Precision and Recall:** For each synthetic tree, vertices of the output skeleton with a corresponding point in the ground truth skeleton within a 0.02m radius are labeled as true positives ( $TP$ ), while vertices without correspondence are labeled as false positives ( $FP$ ). Vertices of the ground truth skeleton without a corresponding point in the output skeleton are labeled as false negatives ( $FN$ ).
- **Occluded Skeleton Ratio (OSR):** To quantify the effectiveness of our approach with regards to predicting occluded branch skeletons correctly, we measure the percentage of true positive vertices that were obtained from the minimum cost paths ( $TP_{occ}$ ) versus the total number of vertices in the constructed skeleton:

$$OSR = \frac{TP_{occ}}{TP + FP} \quad (5)$$

We compare our results to those obtained from a minimum spanning tree (MST) based method [16], [21] as well as a method presented in FTSEM [17]. For the MST-based method, our process of connecting disjoint subgraphs via searching for minimum cost paths in the likelihood map (described in Sec. III-D) is replaced by returning the Euclidean minimum spanning tree built from  $V_{init}$  as the final skeleton. For FTSEM, we use the breakpoint connection method [17] to connect disjoint subgraphs by checking for distance and angle conditions between vectors computed from edges in subgraph pairs. Since the code for FTSEM is not open sourced, we implement it ourselves based on the available details.

The results are summarized in Table I and visualized in Fig. 8. The precision, recall, and OSR of our method averaged over all 120 synthetic trees is 0.98, 0.59 and 0.11,

TABLE I  
EXPERIMENT RESULTS OF TREE SKELETONIZATION  
ON SYNTHETIC TREES.

Tree Type	FD	MST			FTSEM			Ours		
		P	R	OSR	P	R	OSR	P	R	OSR
Oak	1	<b>0.99</b>	0.88	<b>0.05</b>	0.97	0.86	0.03	<b>0.99</b>	<b>0.91</b>	0.04
	2	0.96	0.78	<b>0.08</b>	0.93	0.73	0.04	<b>0.98</b>	<b>0.82</b>	<b>0.08</b>
	3	0.89	0.51	0.14	0.88	0.46	0.11	<b>0.96</b>	<b>0.54</b>	<b>0.15</b>
	4	0.86	0.29	0.17	0.90	0.27	0.16	<b>0.98</b>	<b>0.32</b>	<b>0.19</b>
Apple	1	0.94	0.71	<b>0.14</b>	0.91	0.64	0.07	<b>0.98</b>	<b>0.73</b>	0.12
	2	0.95	0.62	<b>0.16</b>	0.94	0.51	0.05	<b>0.99</b>	<b>0.65</b>	0.13
	3	0.95	0.49	<b>0.14</b>	0.91	0.40	0.05	<b>0.99</b>	<b>0.52</b>	<b>0.14</b>
	4	0.94	0.39	0.15	0.91	0.32	0.06	<b>0.97</b>	<b>0.42</b>	<b>0.16</b>
Walnut	1	0.95	0.70	<b>0.10</b>	0.95	0.65	0.07	<b>0.99</b>	<b>0.72</b>	0.08
	2	0.93	0.64	<b>0.12</b>	0.93	0.56	0.06	<b>0.98</b>	<b>0.66</b>	0.10
	3	0.86	0.48	0.15	0.84	0.39	0.10	<b>0.96</b>	<b>0.49</b>	<b>0.16</b>
	4	0.80	<b>0.34</b>	0.16	0.84	0.29	0.13	<b>0.97</b>	<b>0.34</b>	<b>0.17</b>

\* **FD:** Foliage Density, **P:** Precision, **R:** Recall

\* Each score is averaged over 10 trees belonging to the tree type and foliage density.

respectively. Our method outperforms both baselines in terms of precision and recall as it joins disconnected skeletons only when a path in the likelihood map exists. In contrast, the MST-based method greedily joins all disconnected skeletons. Although this results in higher OSR in less occluded scenes (Table I, foliage density 1 & 2), the relative performance of MST deteriorates with increasing foliage density (Table I, foliage density 3 & 4). As a result, our method outperforms MST in terms of precision, recall and OSR where there is high occlusion.

### B. Experiments on Real Trees

Experiments on real trees were performed on a dataset that consists of 7 apple trees (Fig. 9) collected from the




Fig. 9. Experiment results of seven real apple trees. The first row shows the images of the apple tree. The region bounded by the red dashed square is imaged by the robot with a stereo camera, corresponding to the 3D reconstructed point cloud depicted in the second row. The result of our skeletonization method is shown in the third row with recovered volume by replacing skeleton vertices with spheres of corresponding radius.




Fig. 10. (a) Our hardware setup for collecting real world data at the UMass Cold Spring Orchard, featuring a stereo camera attached to the UR5 robot arm. (b) The robot motion path is shown by the red arrows, where it collects 70 stereo image pairs of each tree canopy at equally distributed waypoints throughout the motion path.

University of Massachusetts Amherst Cold Spring Orchard. For each tree, stereo images were collected from 70 different viewpoints using a flash stereo camera [22] attached to the UR5 robot arm on the Bumblebee system [23] (Fig. 10(a)). The robot was controlled to follow a motion path as depicted in Fig. 10(b), where the viewpoints are set to be at equally distributed waypoints throughout the motion path. Due to the robot's limited workspace, we were able to capture a range of approximately 1.5 meters in height for each tree canopy.

Our results are visualized in Fig. 9. The OSR averaged over all seven trees is 0.14, consistent with the results obtained from the synthetic dataset. As it is difficult to obtain the ground truth skeleton for real trees, we visually compare the extracted skeleton with the 3D reconstructed tree using known camera poses from the robot. Despite considerable occlusion and noise evident in the reconstructed point cloud, the extracted skeleton is topologically correct and shows good correspondence.

The quality of the skeleton obtained from our pipeline is contingent on the sufficiency and correctness of the detected branches. We expect that a control policy to collect images from optimal viewpoints [24] (rather than fixed viewpoints as in our experiments) to perceive sufficient amount of branches will further improve our proposed skeletonization pipeline.

## V. CONCLUSION

Our tree skeletonization method outperforms the baselines in situations with highly occluded canopies by accurately estimating unobserved skeletons. As future work, we plan to improve the algorithm runtime in addition to conducting a more rigorous evaluation on real tree datasets. Other possible directions for future work include next-best viewpoint optimization to increase the information of occluded regions in the tree canopy. We are also interested in estimating the dynamics of the tree in response to external forces to plan for contact interactions. Ultimately, the digitized model of a tree crop in the form of a skeleton presents promising direction to developing safe and robust agricultural robotic manipulation.

## ACKNOWLEDGMENT

This work was supported in part by NSF/USDA-NIFA Cyber Physical Systems 2020-67021-31531, NSF Robust Intelligence 1956163, and NSF/USDA-NIFA AIIRA AI Research Institute 2021-67021-35329. The authors would like to thank Daniel Cooley, Paul O'Connor, Jon Clements, Harry Freeman and Abhisesh Silwal for their help in field data collection at the University of Massachusetts Amherst Cold Spring Orchard, and Hung-Jui Huang for helpful discussions.

## REFERENCES

- [1] H. Huang, S. Wu, D. Cohen-Or, M. Gong, H. Zhang, G. Li, and B. Chen, "L1-medial skeleton of point cloud," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 65–1, 2013.
- [2] J. Cao, A. Tagliasacchi, M. Olson, H. Zhang, and Z. Su, "Point cloud skeletons via laplacian based contraction," in *2010 Shape Modeling International Conference*, pp. 187–197, IEEE, 2010.
- [3] A. Tagliasacchi, H. Zhang, and D. Cohen-Or, "Curve skeleton extraction from incomplete point cloud," in *ACM SIGGRAPH 2009 papers*, pp. 1–9, 2009.
- [4] F. Yandun, A. Silwal, and G. Kantor, "Visual 3d reconstruction and dynamic simulation of fruit trees for robotic manipulation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 54–55, 2020.
- [5] A. You, H. Kolano, N. Parayil, C. Grimm, and J. R. Davidson, "Precision fruit tree pruning using a learned hybrid vision/interaction controller," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2280–2286, IEEE, 2022.
- [6] D. Sepúlveda, R. Fernández, E. Navas, M. Armada, and P. González-De-Santos, "Robotic aubergine harvesting using dual-arm manipulation," *IEEE Access*, vol. 8, pp. 121889–121904, 2020.
- [7] C. Zhang, G. Yang, Y. Jiang, B. Xu, X. Li, Y. Zhu, L. Lei, R. Chen, Z. Dong, and H. Yang, "Apple tree branch information extraction from terrestrial laser scanning and backpack-lidar," *Remote Sensing*, vol. 12, no. 21, p. 3592, 2020.
- [8] D. Zhang, N. Xie, S. Liang, and J. Jia, "3d tree skeletonization from multiple images based on pyrllk optical flow," *Pattern Recognition Letters*, vol. 76, pp. 49–58, 2016. Special Issue on Skeletonization and its Application.
- [9] L. Fu, J. Liu, J. Zhou, M. Zhang, and Y. Lin, "Tree skeletonization for raw point cloud exploiting cylindrical shape prior," *IEEE Access*, vol. 8, pp. 27327–27341, 2020.
- [10] A. You, C. Grimm, A. Silwal, and J. R. Davidson, "Semantics-guided skeletonization of upright fruiting offshoot trees for robotic pruning," *Computers and Electronics in Agriculture*, vol. 192, p. 106622, 2022.
- [11] S. Delagrangé, C. Jauvin, and P. Rochon, "Pypetree: A tool for reconstructing tree perennial tissues from point clouds," *Sensors*, vol. 14, no. 3, pp. 4271–4289, 2014.
- [12] A. Bucksch, R. Lindenbergh, and M. Menenti, "Skeltre," *The Visual Computer*, vol. 26, no. 10, pp. 1283–1300, 2010.
- [13] J. Hackenberg, C. Morhart, J. Sheppard, H. Spiecker, and M. Disney, "Highly accurate tree models derived from terrestrial laser scan data: A method description," *Forests*, vol. 5, no. 5, pp. 1069–1105, 2014.
- [14] J. Mei, L. Zhang, S. Wu, Z. Wang, and L. Zhang, "3d tree modeling from incomplete point clouds via optimization and l1-mst," *International Journal of Geographical Information Science*, vol. 31, no. 5, pp. 999–1021, 2017.
- [15] L. Gao, D. Zhang, N. Li, and L. Chen, "Force field driven skeleton extraction method for point cloud trees," *Earth Science Informatics*, vol. 12, no. 2, pp. 161–171, 2019.
- [16] Y. Livny, F. Yan, M. Olson, B. Chen, H. Zhang, and J. El-Sana, "Automatic reconstruction of tree skeletal structures from point clouds," in *ACM SIGGRAPH Asia 2010 papers*, pp. 1–8, 2010.
- [17] J. Sun, P. Wang, R. Li, M. Zhou, and Y. Wu, "Fast tree skeleton extraction using voxel thinning based on tree point cloud," *Remote Sensing*, vol. 14, no. 11, p. 2558, 2022.
- [18] M. Bremer, M. Rutzinger, and V. Wichmann, "Derivation of tree skeletons and error assessment using lidar point cloud data of varying quality," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 80, pp. 39–50, 2013.
- [19] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- [20] L. Piegl and W. Tiller, *The NURBS book*. Springer Science & Business Media, 1996.
- [21] B. Gorte and N. Pfeifer, "Structuring laser-scanned trees using 3d mathematical morphology," *International Archives of Photogrammetry and Remote Sensing*, vol. 35, no. B5, pp. 929–933, 2004.
- [22] A. Silwal, T. Parhar, F. Yandun, H. Baweja, and G. Kantor, "A robust illumination-invariant camera system for agricultural applications," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3292–3298, 2021.
- [23] A. Silwal, F. Yandun, A. Nellithimaru, T. Bates, and G. Kantor, "Bumblebee: A path towards fully autonomous robotic vine pruning," *arXiv preprint arXiv:2112.00291*, 2021.
- [24] X. Zeng, T. Zaenker, and M. Bennewitz, "Deep reinforcement learning for next-best-view planning in agricultural applications," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2323–2329, IEEE, 2022.