

Ground-Density Clustering for Approximate Agricultural Field Segmentation

Henry J. Nelson and Nikolaos Papanikolopoulos

{nels8279, papan001}@umn.edu

Department of Computer Science and Engineering, University of Minnesota

Abstract—Instance and semantic segmentation form the backbone of robotic perception and are crucial to many tasks. While most research in the area focuses on improving segmentation quality metrics, there are plenty of applications where approximate methods are adequate as long as they are fast, especially in applications with large amounts of data like precision agriculture. In order to apply the recent successes of machine learning and computer vision on a large scale using robotics, efficient and general algorithms must be designed to intelligently split point clouds into small, yet actionable, portions that can then be processed by more complex algorithms. In this paper, we capitalize on a similarity between the current state-of-the-art for roughly segmenting corn plants and a commonly used density-based clustering algorithm, Quickshift. Exploiting this similarity we propose a novel algorithm, Ground-Density Quickshift++, with the goal of producing a general and scalable field segmentation algorithm that segments individual plants and their stems. This algorithm produces quantitatively better results than the current state-of-the-art on both plant separation and stem segmentation while being less sensitive to input parameters and maintaining the same algorithmic time complexity. When incorporated into field-scale phenotyping systems, the proposed algorithm should work as a drop-in replacement that can greatly improve the accuracy of results while ensuring that performance and scalability remain undiminished.

I. INTRODUCTION

While the demand for food increases and the need for environmentally conscious farming practices is clearer than ever before, precision agriculture and agricultural robotics promise the ability to reduce input costs, while simultaneously maximizing yield and reducing environmental impact. But to live up to these promises, phenotype data must be collected for individual plants on an extremely large scale [1]. The current solution, which is to estimate the phenotype of an entire field from a manually measured sample, is prone to sampling bias and does not scale. Requiring human intervention to produce data slows the rate at which data can be collected and makes every data point costly to obtain in terms of both time and money. The remote sensing and agricultural robotics communities have been working to automate this process, but obtaining accurate measurements of individual plants in an unstructured environment is a multifaceted problem with many challenging intermediate steps. Currently proposed solutions utilize specialized sensors and computationally expensive algorithms that make these solutions unable to perform at the scale of entire fields. But instead of performing all operations on a large scale, it may be possible to collect data on a large scale




Fig. 1: The result of Ground Density Quickshift++ when run on a point cloud of a corn field. A total of 167 clusters are found with only one case of a cluster containing more than one plant (shown on the right of Fig. 6). There are 136 plants in the input point cloud meaning that some clusters, like those visible between rows in the center of this image, contain no plants.

using a robotic platform and then split the data into small, yet semantically meaningful, units that can be processed individually. This would allow data to be captured at a large scale and then broken into individual plants to be processed by a more computationally expensive phenotyping system, thus avoiding the need to feed large amounts of data into phenotyping systems while simultaneously simplifying the task they need to accomplish. It is therefore necessary to find a way of splitting a field of plants into individual plants in a computationally inexpensive way.

In this paper, we present and evaluate an algorithm developed to roughly segment individual plants and their stems over an entire field. This algorithm is defined as a modification of a general clustering technique so that its computational complexity is low and well-defined. The use of common data clustering techniques as an algorithmic basis constrains the new algorithm to be relatively data and species-agnostic while maintaining a known behavior when scaling up to large input sizes. This algorithm allows data collection and phenotype measurement to happen at different scales so the computationally expensive phenotyping algorithms can operate on small portions of the point cloud containing individual plants in order to maintain a lower runtime.

Because the plant separation problem is just a small part of what is to be a long and complex system for phenotype measurement, the required inputs and outputs of each algorithm are kept as general as possible. The algorithm takes in a point cloud of a group of plants and produces a partition of the input points into sets *roughly corresponding* to individual

plants. It is important to note that since this step is meant as a pre-processing step to produce clusters that will be processed further, the partitioning of plants does not have to be precise and a rough approximation is adequate as long as the algorithm scales. The input point clouds are required to be rotated such that the gravity vector of the world frame is pointing in the $-\vec{z}$ direction (up is $+\vec{z}$ direction), but the coordinate system is otherwise unconstrained. They do not require row, color, or radiometric information. Only the three spatial dimensions are necessary for each point.

The algorithm defined in this work is more general than any other method of separating plants and their constituent parts in point clouds of crop fields. Its independence from a coordinate system, training data, row information, and radiometric or color information make it the most widely applicable algorithm available for the task of approximate plant segmentation. It does not require a specific sensor like lidar or an RGB camera and does not require any training data like common machine learning approaches. One is even unit agnostic, avoiding the need for data to be scaled into known units as is commonly required for reconstructions created via structure-from-motion. This algorithm is not the final solution to the problem, but outperforms the current state-of-the-art for approximate stem and plant segmentation by a large margin [2], is scalable enough to be used on large-scale data, general enough to be a starting point for similar applications, and simple enough to be adapted to specific problems. The algorithm proposed in this work has already been found useful on related agricultural problems [3] and similar approaches have found success even outside the agricultural domain [4].

II. RELATED WORK

For segmenting a field into individual plants, specialized sensors like lidar have proven useful [5]–[8], but lidar sensors remain prohibitively expensive when compared to more commonly available sensors like RGB cameras, which are already commonly used for crop surveying. In [5], radiometric data from the lidar sensor is directly utilized, meaning the algorithm is tied to a single sensing modality. In contrast, [6] and [7] do not use radiometric information so they can presumably be implemented with different sensors, but they instead utilize deep learning methods that require large amounts of training data and are unable to operate at an adequate resolution for field scale point clouds due to the use of voxel representations. Of the works using lidar, the work of [8] is the most general as they use lidar only to obtain the point cloud and then use a region-growing clustering method that incorporates row and plant spacing parameters to separate plants. The resulting plant detection algorithm performs well on simulated data but the performance deteriorates on data from real crop fields, achieving an average detection rate of only 60%.

While using lidar produces a more accurate and complete reconstruction of the environment, it has not been the only approach to crop reconstruction. Some groups have instead opted to use stereo vision or structure-from-motion as they

do not require expensive sensors. However, due to the increased noise and non-uniform sampling imposed by these reconstruction techniques, the approach to the problem is usually slightly different. The work of [9] only operates on data collected in laboratory conditions, [10] uses manually separated plant data, and [11] skips the plant level altogether and just segments individual leaves. In [12] they claim the ability to segment well-separated plants from field data, but the method of separating plants from one another is not detailed enough to be reproducible. The workflow proposed in [2] uses structure-from-motion and then a skeletonization procedure before finally clustering the point cloud to achieve sparse detection and sampled segmentations of plants in a large scale point cloud. This work achieves state-of-the-art results with scalable performance but does so on a skeletonized representation of the point cloud meaning the resulting segmentation is on a more coarse scale than the original point cloud. The algorithm used is easily applicable to the original point cloud though if the skeletonization is not run and the algorithm is performed on the entire point cloud rather than on a sampling, the complexity is still only $\mathcal{O}(n \log n)$ which matches the complexity of the algorithms proposed in this paper and achieves reasonable results. In fact, when extending the RAIN algorithm defined in [2] to run without randomization, a similarity with the Quickshift algorithm [13] was noticed and became the basis of this paper.

Quickshift [13] is a density-based, mode-finding, clustering algorithm. It is similar to Mean-Shift [14] in that points are clustered with those that are drawn to the same mode, but it is not iterative and only evaluates the density function $\mathcal{O}(n)$ times. This makes Quickshift very efficient as its complexity is dominated by spatial neighborhood queries which can be performed quickly with the help of spatial data structures like a k D-Tree. With this paper we sought to exploit the similarity between RAIN and Quickshift to define a more accurate and generalized plant clustering algorithm. In this way, we switch from using algorithms made specifically for a single application to using a slightly modified version of a more general and well-vetted clustering algorithm.

III. ALGORITHMS

Noticing how RAIN [2] takes advantage of point cloud topology and approximates common density based clustering methods, we define a few similar agglomerative clustering algorithms using Quickshift [13] as a template. The first of these is a non-randomized adaptation of the RAIN algorithm to be used for comparison while the others are a modification of Quickshift and Quickshift++ [15] respectively.

A. Non-random RAIN

The first algorithm we define is a non-randomized version of the RAIN algorithm from [2]. The purpose is that RAIN is the current state-of-the-art in this domain but it only partitions a sampling of the points and therefore does not directly match the formulation used in this work. So with this algorithm, we remove the random sampling portion of




Fig. 2: A visualization of a density function on a point cloud of plants projected into the ground plane. The red regions (areas of high density) are easily separable and identify individual plant stems. This is due to the vertical structure of the stem causing all points along the stem to project to the same region of the ground plane.

the algorithm and run the procedure on all points in the point cloud. In this way, it is directly comparable to the other algorithms in this work while maintaining its state-of-the-art performance. The algorithm takes as input a neighborhood distance d and a point cloud $P = \{\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n : \vec{p}_i \in \mathbb{R}^3\}$. For each point $\vec{p} \in P$, we assign \vec{p} to the cluster of the point in $N_d(\vec{p})$ with the minimum z coordinate where $N_d(\vec{p}) = \{\vec{n} \in P : \|\vec{p} - \vec{n}\| < d\}$, the neighborhood of points within distance d of the point \vec{p} . Assuming this neighborhood can be computed with a spatial data structure like a kD -tree in approximately $\mathcal{O}(\log n)$ time, this means the entire algorithm operates in $\mathcal{O}(n \log n)$. This adaptation of the algorithm leaves it with only one parameter, d , which is a neighborhood distance threshold in 3D space.

In addition to separating plants, RAIN was also been shown to enable rough semantic segmentations by separating stems and leaves [2]. By iteratively running the RAIN algorithm multiple times and restricting sampled points to be around cluster centers from the previous iterations, the central portion of clusters can be extracted. This central portion usually represents the stem of each cluster’s plant. This variant of the RAIN algorithm for stem segmentation requires one extra parameter governing which points are close enough to a previous cluster center to be included in the following iteration. This iteration happens a constant number of times keeping the algorithmic complexity the same.

B. Ground Density Quickshift

Instead of utilizing the vertical structure of the plant as done in RAIN, it was noted that viewing the structure of the plant in the horizontal plane gives a very clear picture of the separation between plants and takes advantage of the horizontal density common to many vertically structured plants (corn, trees, hops) to cluster points that are close in the projected space. This structure can be seen in Fig. 2. This is conducive to clustering points that are close in the projected space but may not be close in 3D, such as two points on opposite ends of the plant stem or trunk. To take advantage of the observable structure in the ground plane we can project all points into the ground plane where we can then perform standard Quickshift in 2D. For this algorithm we utilize the density function defined in [15] as it can be performed

with the neighborhood queries which must be computed for Quickshift anyway, allowing a runtime of $\mathcal{O}(n \log n)$. The density function has a single integer parameter k determining the kernel size. This gives *GD Quickshift* a total of two parameters, d and k .

C. Ground Density Quickshift++

Projecting the data into the ground plane before clustering easily separates vertical structures but will also place leaves in the same cluster if they overlap in the 2D space even if they are far apart in 3D. To combat this we could instead use the densities determined in 2D and cluster the points in 3D, but this creates many small clusters along the length of the stem rather than keeping them all in the same cluster. As a workaround, we can instead use Quickshift++ [15] and cluster the densest regions of nearby points in the 2D density space as a set of seed clusters. We can then cluster the remaining points in 3D using Quickshift to ensure points are only in the same cluster if they are close in 3D or were in an extremely dense region in the 2D projection (probably a stem). Due to the initialization, there is no need for a distance parameter making this algorithm completely agnostic to scale. Instead, Quickshift++ uses a parameter $\beta \in [0, 1]$ to identify the densest regions of points as cluster seeds. Along with the k parameter used for the density function, *GD Quickshift++* has a total of two parameters. The initialization of clusters runs in $\mathcal{O}(n \log n)$ so the final algorithm is also $\mathcal{O}(n \log n)$.

Similarly to RAIN, we can also define a variant of the algorithm to segment stems from the point cloud. In fact, the initial clusters taken from the densest regions of the point cloud when projected into the ground plane are exactly this. These cluster seeds can be saved to be used as a rough stem segmentation before also being used as the seeds of the subsequent 3D clustering step. This does not require any additional work and thus maintains the same algorithmic complexity while also producing approximate stem segmentations.

IV. DATA

To evaluate the algorithms from Section III, both synthetic data and real-world sensor data are used. The corn field data and implementations of the algorithms have been made publicly available¹ to facilitate the replication of results and continued development by the community. The non-corn data used in this work are already publicly available [16].

Synthetic data was created by taking the individual plant reconstructions provided in [17] and organizing them into a large-scale field of plants. This allows us to have a known ground truth partitioning of the point cloud with which to quantitatively compare. Two large-scale point clouds are synthesized for the plant separation task, one containing 100 plants and the other containing 400 plants. The ground truth partitions are encoded in the point clouds using color information. Additionally, a second point cloud of 400 plants was

¹<https://github.com/hennels/CropPreClustering>




Fig. 3: Qualitative results for a portion of the St. Paul V3 data. Each column is a different algorithm and each row shows a different value of the neighborhood distance threshold d . The equivalent image for Quickshift++ is shown on the left of Fig. 4.

produced to evaluate each algorithm’s semantic segmentation of plant stems, a selection of which can be seen on the right side of Fig. 4.

A small dataset of reconstructions from production corn fields was also compiled. A UAV with an attached RGB camera flew at a low altitude over corn fields collecting video. Afterward, the video was used to build a 3D reconstruction of the corn plants using Pix4D to perform structure-from-motion. The dense central region of each reconstruction was then extracted and transformed such that the plant stems were parallel with the \bar{z} -axis. The data was collected in three different fields over two different growing seasons with corn plants spanning stages from V3 to V6. Qualitative results from only one point cloud are shown in this paper (in Figures 1, 3, and 4) but the reader is encouraged to test the algorithms on the other provided point clouds to see additional results.

V. EVALUATION

In order to evaluate the algorithms proposed in this paper we use both quantitative and qualitative methods. For the synthetic data with ground truth instance partitions, we are able to directly compare the results of the algorithms with the ground truth after finding a bipartite correspondence between predicted and ground truth clusters. To compare partitions the intersection over union (IoU) is used and a bipartite matching between predicted and ground truth clusters is computed such that the sum of IoUs is maximized. This makes sure each estimated cluster is only compared against a single ground truth cluster. For ground truth data with semantic stem labels, the Dice score is computed between the segmented set and the ground truth points of the stem class.

For the real-world data where ground truth is not available, we show qualitative results for each algorithm using a variety of parameters. To evaluate, we run each algorithm

Algorithm	Clusters	mean IoU	median IoU
Ground Truth	400	-	-
Non-random RAIN	474	62.1%	68.4%
GD Quickshift	397	81.0%	81.7%
GD Quickshift++	401	84.3%	85.1%

TABLE I: Quantitative results on 10to29.1M.ply for each algorithm. A single IoU value is computed between each predicted and ground truth cluster that are matched by the bipartite matching. The mean and median values of these matched pairs are reported.

on the same point cloud, with the same parameters, and view the differences. Since each algorithm requires different parameters, we notice that the algorithms are most sensitive to the neighborhood distance threshold d , so this is the parameter we vary. For k , a few values were tried and one with reasonable performance on all point clouds ($k = 1200$) is shown. The GD Quickshift++ algorithm was found to give good results for any value of $\beta \in [0.1, 0.4]$ so a value of $\beta = 0.3$ is used for all work shown in this paper. For the best-performing algorithm, we are able to compare manually verifiable metrics on real-world data such as the ground truth number of plants in the point cloud vs. the total number of clusters and the number of clusters containing more than one plant because the errors are few enough to be manually counted.

Since these are pre-clustering algorithms, it is important to note that the necessary performance is dependent on what the subsequent processing algorithms can handle. Ideally, a plant pre-clustering algorithm should be able to separate plants and keep all plants whole in the process, but these goals are often in opposition to one another and it is difficult to accomplish both well. So depending on the subsequent processing algorithm, an over or under-fragmented segmentation may be preferable and the parameters could be adjusted to produce such a partition.

A. Quantitative Evaluation

For quantitative evaluation of the instance segmentation performance, each algorithm was run on the synthetic data with a range of parameters. The results in Tables I and II show the best results for each algorithm that produced a number of clusters within 20% of the ground truth value on the two point clouds. Due to the bipartite matching between estimated and ground truth clusters, the IoU values can be significantly biased if the number of estimated clusters differs significantly from the number of ground truth clusters which is the reason for the 20% bound. In both Tables I and II Ground Density Quickshift++ outperforms the other algorithms on all metrics and thus this is the algorithm we

Algorithm	Clusters	mean IoU	median IoU
Ground Truth	100	-	-
Non-random RAIN	102	54.9%	65.8%
GD Quickshift	116	81.9%	82.7%
GD Quickshift++	100	85.7%	87.1%

TABLE II: Quantitative results on 0to9.1M.ply for each algorithm. A single IoU value is computed between each predicted and ground truth cluster that are matched by the bipartite matching. The mean and median values of these matched pairs are reported.

Dataset	Total Plants	Total clusters	Multi-plant Clusters
St. Paul V3	136	144	4
St. Paul V6	103	107	6
Dji V4	90	92	3

TABLE III: Quantitative results for GD Quickshift++ running on three different point clouds. Note that the parameters of GD Quickshift++ were not tuned for any one of these point clouds, the values of $k = 1200$ and $\beta = 0.3$ were chosen as parameters that work reasonably well for all of them. Improvement for each would be expected with parameter tuning.

also evaluate quantitatively on the semantic segmentation task and real-world data.

For the stem segmentation task, an extensive parameter sweep was performed using the stem segmentation variants of Nonrandom RAIN and Ground-Density Quickshift++ on the synthetic point cloud with stem labels. The Dice score between the predicted and ground truth point sets were computed for each parameter combination and the maximum for each algorithm is reported here to show the best possible outcome of each algorithm. On the stem segmentation task, Nonrandom RAIN achieved a Dice score of 0.62 while **GD Quickshift++ achieved a Dice score of 0.88**, meaning that GD Quickshift++ produces a much higher quality stem segmentation with the same computational complexity. The results can be seen visually on the right side of Fig. 4. We can then turn our attention to the performance of these algorithms on real-world data.

After applying Ground Density Quickshift++ to each of our real-world point clouds the number of plants and the number of clusters containing more than one plant were manually counted in each. The number of clusters with more than one plant is a metric measuring the cases where the algorithm was unable to separate plants. All of these occurrences were clusters with only two plants, most of which were tillers [18] or double planting events as shown in Fig. 6. Double planting events are events where two seeds are planted in the same location because of a malfunction in the planting machinery. The number of clusters found minus the total number of plants in the point cloud is a measure of how many extraneous clusters the algorithm produces. Ideally, the number of multi-plant clusters and the number of extra clusters should both be zero. While we considered any cluster not containing a plant of the cultivated crop extraneous, it was noted that many of these clusters were on weeds in the field, so they may be desirable for applications such as weed monitoring or planning herbicide application. The results are shown in Table III. From these values, we can see that the total number of clusters produced by Ground Density Quickshift++ is only slightly higher than the total number of plants and the number of multi-plant clusters is relatively small for all three point clouds.

B. Qualitative Evaluation

In Figures 3 and 4 qualitative results on the St. Paul V3 dataset are shown. From these images, it is clear that non-randomized RAIN (left column of Fig. 3) is unable to completely separate plants with any of the sampled parameters and often produces more than one cluster on each plant.




Fig. 4: **Left:** Qualitative results for Quickshift++ on a portion of the St. Paul V3 data. It is directly comparable to the results in Fig. 3 but does not use a neighborhood distance threshold. See Fig. 1 for a view of the whole field. **Right:** Example point cloud showing GD Quickshift++ stem segmentation. Yellow: True positive stem points, Blue: True negative leaf points, Magenta: False positive stem, Green: False negative stem.

The behavior changes quickly from over-fragmentation to under-fragmentation as the neighborhood distance threshold increases. Ground Density Quickshift (right column of Fig. 3) gives comparatively much cleaner segmentations and seems less sensitive to the neighborhood distance threshold as the over and under-fragmentation are well balanced for a majority of the d input values. The Ground Density Quickshift++ result, which is shown on the left side of Fig. 4 as it does not utilize a neighborhood distance threshold, is the cleanest of any of the algorithms tested and is very insensitive to input parameters. Each plant in the Ground Density Quickshift++ result is part of just a single cluster while the ground points are placed in relatively few clusters. The performance of Ground Density Quickshift++ on a larger portion of the same dataset can be seen in Fig. 1.

While this work focuses on corn segmentation, the algorithm shows promise for a variety of crops. As mentioned previously, during experiments on data from corn fields a number of weeds were found to also be well segmented by the algorithm. To investigate whether the algorithms are effective on other crop types, Ground Density Quickshift++ was run on a publicly accessible point cloud of Fuji apple trees [16]. The result is visible in Fig. 5. The algorithm was able to produce a reasonable segmentation of the apple trees in the point cloud, making only a few small errors. Two of the apple trees were unable to be separated and one side of a tree that was particularly dense with apples was counted as a separate tree, making the total number of clusters equal to the number of trees captured in the reconstruction. While the apple tree dataset is a small example, it and the weeds that were segmented from the corn datasets do suggest that Ground Density Quickshift++ is extensible to some other crop domains with little to no modification.




Fig. 5: The result of running Quickshift++ on structure from motion data of Fuji apple trees from [16]. Clusters mostly represent individual trees. Minor errors in segmentation balance out so that the total number of trees detected is correct.

VI. DISCUSSION

Although all the algorithms presented in this work have a similar structure, they have very different behaviors as is seen in Section V. The non-random RAIN algorithm does not group nearby points well because its goal is to group points with their lowest neighbor within some distance. This does not ensure that points cluster with their immediate neighbors and often results in a periodic clustering pattern as is shown on the left of Fig. 6. This shows the need for points to be clustered with neighbors as long as they are part of the same structure. This motivates the use of the density function used in the Ground Density algorithms. Both of the ground density function-based approaches work much better in general and produce extremely clean clusters that mostly represent individual plants.

The main difference between Ground Density Quickshift and Ground Density Quickshift++ is the presence or absence of the neighborhood distance threshold. The neighborhood parameter governs how close two points need to be in order to be clustered together. This means it will have to be less than the distance between the nearest two plants if they are to be separated, but this is difficult in some situations like the double planting events shown in Fig. 6. The absence of the threshold in Ground Density Quickshift lets the threshold differ across the point cloud. Plants are not evenly distributed in a field, so typically the threshold would need to be lower in the row of plants and larger in the gap between rows to minimize the number of extraneous clusters. Requiring a constant threshold ensures that either plants will not be separated or there will be a large number of extraneous clusters. One of these alternatives may be permissible because they are able to be handled by subsequent algorithms but this deficiency adds algorithmic complexity later in the processing pipeline. Ground Density Quickshift++ does not require a constant threshold so it is better able to adapt to these different plant densities. This explains its good performance on the areas of different density seen in Figs. 4 and 1 where it has many clusters close together in the row of plants but very few in the gap between plant rows.

Without the explicit neighborhood distance threshold, Ground Density Quickshift++ is independent of the scale of the point cloud meaning a point cloud produced via structure-




Fig. 6: **Left:** A common failure mode of non-random RAIN. Since RAIN is always looking for the lowest point, it can produce layered clusters on the same plant. **Right:** Two examples of double planting events. Plants from two different seeds that are planted in the same place. The example on the left shows an instance of close plants that the algorithm was unable to separate while the instance on the right is correctly separated.

from-motion does not need to be scaled to world units before processing. The other three algorithms all require the neighborhood distance threshold and thus the choice of d is dependent on the units of the point cloud. So to use consistent parameters, point clouds would need to be scaled to world units or produced using a method with a known scale like lidar. Even though Ground Density Quickshift++ is independent of scale, it is dependent on the sampling density of the point cloud due to its density function. If scale or point cloud density is easier to control, this difference may inform which algorithm may be of use in a given application.

VII. CONCLUSIONS

In this paper, we proposed and evaluated a range of scalable clustering algorithms designed for use in entire fields. The best of these, Ground Density Quickshift++, performs better than the current state-of-the-art on both approximate stem and individual plant segmentation. While their performance differs, all algorithms in this work have a time complexity of approximately $\mathcal{O}(n \log n)$. We also provided multiple synthetic datasets with ground truth instances or semantic segmentations and a real-world dataset on which to test the proposed algorithms and foster further development in this domain.

This work enables the development of a high throughput plant phenotyping system for agricultural robotics by providing algorithms to act as a first step to break a reconstruction into manageable parts representing individual plants; a task that any robotic system will have to accomplish in order to apply techniques from precision agriculture. This allows data capture to occur at field scale without requiring the subsequent phenotyping algorithms to operate at the same large scale, allowing them to be more complex and specialized. More specialized algorithms will need to be developed to operate on the resulting clusters, but many techniques have already shown promising results in such niche contexts [3], [19]. The main contribution of this work is not only that the proposed algorithms outperform the current state-of-the-art method with the same time complexity, but also the generality and thus applicability of the proposed algorithms which have already found use in agricultural contexts [3], [20], and other domains [4]. As long as plants have a single vertical structure (like corn, fruit trees, or hops), one of the Ground Density algorithms should be able to find and segment them.

VIII. ACKNOWLEDGEMENTS

We would like to thank all members of the Center for Distributed Robotics at the University of Minnesota for help with data collection and fruitful discussion. A special thanks goes to Dimitris Zermas and Sentera, Inc. for processing imagery. USDA/NIFA has supported this work through the grants 2020-67021-30755 and 2023-67021-39829. This work was partially supported by the National Science Foundation through grants #CNS-1439728, #CNS-1544887, and #CNS-1939033.

REFERENCES

- [1] J. L. Araus and J. E. Cairns, "Field high-throughput phenotyping: the new crop breeding frontier," *Trends in plant science*, vol. 19, no. 1, pp. 52–61, 2014.
- [2] D. Zermas, V. Morellas, D. Mulla, and N. Papanikolopoulos, "Extracting phenotypic characteristics of corn crops through the use of reconstructed 3d models," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 8247–8254.
- [3] D. Wang, Z. Song, T. Miao, C. Zhu, X. Yang, T. Yang, Y. Zhou, H. Den, and T. Xu, "Dfsp: A fast and automatic distance field-based stem-leaf segmentation pipeline for point cloud of maize shoot," *Frontiers in Plant Science*, vol. 14, p. 1109314, 2023.
- [4] Y. He, L. Jin, B. Guo, Z. Huo, H. Wang, and Q. Jin, "Density-based road segmentation algorithm for point cloud collected by roadside lidar," *Automotive Innovation*, vol. 6, no. 1, pp. 116–130, 2023.
- [5] B. Höfle, "Radiometric correction of terrestrial lidar point cloud data for individual maize plant detection," *IEEE Geoscience and Remote Sensing Letters*, vol. 11, no. 1, pp. 94–98, 2014.
- [6] S. Jin, Y. Su, S. Gao, F. Wu, T. Hu, J. Liu, W. Li, D. Wang, S. Chen, Y. Jiang, S. Pang, and Q. Guo, "Deep learning: Individual maize segmentation from terrestrial lidar data using faster r-cnn and regional growth algorithms," *Frontiers in Plant Science*, vol. 9, p. 866, 2018. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fpls.2018.00866>
- [7] S. Jin, Y. Su, S. Gao, F. Wu, Q. Ma, K. Xu, Q. Ma, T. Hu, J. Liu, S. Pang, H. Guan, J. Zhang, and Q. Guo, "Separating the structural components of maize for field phenotyping using terrestrial lidar data and deep convolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 4, pp. 2644–2658, 2020.
- [8] U. Weiss and P. Biber, "Plant detection and mapping for agricultural robots using a 3d lidar sensor," *Robotics and autonomous systems*, vol. 59, no. 5, pp. 265–273, 2011.
- [9] F. Golbach, G. Kootstra, S. Damjanovic, G. Otten, and R. van de Zedde, "Validation of plant part measurements using a 3d reconstruction method suitable for high-throughput seedling phenotyping," *Machine Vision and Applications*, vol. 27, no. 5, pp. 663–680, 2016.
- [10] P. Sodhi, S. Vijayarangan, and D. Wettergreen, "In-field segmentation and identification of plant structures using 3d imaging," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 5180–5187.
- [11] B. Biskup, H. Scharf, U. Schurr, and U. Rascher, "A stereo imaging system for measuring structural parameters of plant canopies," *Plant, cell & environment*, vol. 30, no. 10, pp. 1299–1308, 2007.
- [12] S. Jay, G. Rabatel, X. Hadoux, D. Moura, and N. Gorretta, "In-field crop row phenotyping from 3d modeling performed using structure from motion," *Computers and Electronics in Agriculture*, vol. 110, pp. 70–77, 2015.
- [13] A. Vedaldi and S. Soatto, "Quick shift and kernel methods for mode seeking," in *European conference on computer vision*. Springer, 2008, pp. 705–718.
- [14] Y. Cheng, "Mean shift, mode seeking, and clustering," *IEEE transactions on pattern analysis and machine intelligence*, vol. 17, no. 8, pp. 790–799, 1995.
- [15] H. Jiang, J. Jang, and S. Kpotufe, "Quickshift++: Provably good initializations for sample-based mean shift," in *International Conference on Machine Learning*. PMLR, 2018, pp. 2294–2303.
- [16] J. Gené-Mola, R. Sanz-Cortiella, J. R. Rosell-Polo, J.-R. Morros, J. Ruiz-Hidalgo, V. Vilaplana, and E. Gregorio, "Fruit detection and 3d location using instance segmentation neural networks and structure-from-motion photogrammetry," *Computers and Electronics in Agriculture*, vol. 169, p. 105165, 2020.
- [17] H. J. Nelson and N. Papanikolopoulos, "Learning continuous object representations from point cloud data," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 2446–2451.
- [18] R. L. Veenstra, C. D. Messina, D. Berning, L. A. Haag, P. Carter, T. J. Hefley, P. V. Prasad, and I. A. Ciampitti, "Effect of tillers on corn yield: Exploring trait plasticity potential in unpredictable environments," *Crop Science*, vol. 61, no. 5, pp. 3660–3674, 2021.
- [19] T. Miao, C. Zhu, T. Xu, T. Yang, N. Li, Y. Zhou, and H. Deng, "Automatic stem-leaf segmentation of maize shoots using three-dimensional point cloud," *Computers and Electronics in Agriculture*, vol. 187, p. 106310, 2021.
- [20] G. Roggiolani, F. Magistri, T. Guadagnino, J. Behley, and C. Stachniss, "Unsupervised pre-training for 3d leaf instance segmentation," *IEEE Robotics and Automation Letters*, 2023.