2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

# Robust Plant Localization and Phenotyping in Dense 3D Point Clouds for Precision Agriculture

Henry J. Nelson, Christopher E. Smith*, Athanasios Bacharis and Nikolaos P. Papanikolopoulos
{nels8279, bacha035, papan001}@umn.edu, csmith16@lssu.edu
*Department of Computer Science and Engineering, University of Minnesota*
**School of Computer Science and Mathematics, Lake Superior State University*

***Abstract*— The determination of a crop's growth-stage is critical information for precision agriculture. Estimates of the growth-stage are used to guide irrigation and the application of agrochemicals. Of particular importance is the use of fertilizers, however, growth-stage estimates may also suggest further investigation of potential crop infections and infestations. Traditionally, the growth-stage is based upon a manual random sample of a very small number of plants that are then analyzed to produce an estimate for the entire crop (up to thousands of acres). In order to increase the sample size (and thus accuracy) and to enable precision agriculture to address non-uniform crop development across a field, we present an analysis methodology that facilitates the automated growth-stage analysis of dense point clouds that are derived from drone imagery. Our method utilizes a standard camera drone and does not use specialized sensors or geo-spatial tagging. We propose a multi-stage unsupervised method, which provides information about the individual plant locations in a field plot with a high probability. The method also produces a measure of individual plant heights, which along with their location are critical for later growth-stage estimation and necessary for robotic precision application. We confirm our method's efficacy with experimental results on corn fields in Minnesota.**

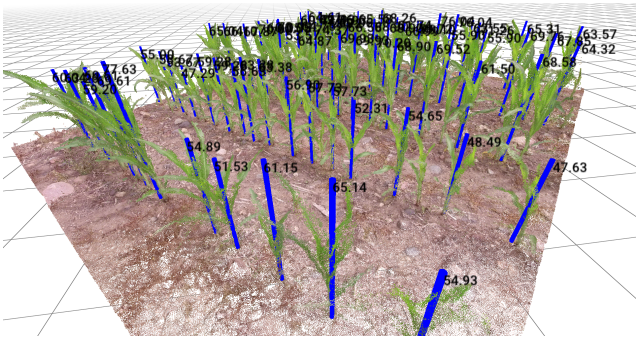


**Fig. 1:** Results of the proposed method on a point cloud of row crops (Cloud 1) showing the detected plant locations in blue and the measured plant heights in *cm*.

## I. Introduction

One way to reduce the impact of agriculture on our environment is to build and deploy robotic systems for precision agriculture. Intelligent systems can collect crop data, use that data to determine plant health, and plan a mitigation and/or remediation approach that targets only the affected plants. Once the plants are identified and diagnosed, robotic systems can carry out the proper treatment autonomously with plant-level precision rather than mass treatment of the entire crop [1]. The benefits of such systems are broad.

First, the cost of irrigation and agrochemicals will be reduced, thus increasing the profit margin for the farmer [1]. Second, the amount of chemical runoff into waterways, aquifers, and other bodies of water will be reduced [1]. This reduces algae blooms in lakes and streams, reduces the creation of so-called "dead-zones" in large bodies of water, and reduces the contamination of aquifers and other sources of drinking water, particularly those in rural areas. Nitrogen and phosphorous from fertilizers can lead to eutrophication (phytoplankton and algae blooms) in watersheds, this in turn can result in low oxygen levels that create "dead zones." Nitrogen also easily enters groundwater, thereby affecting wells used for drinking water in rural farmland. Nitrates in drinking water are known to affect the health of people, particularly infants [2] [3]. The third benefit is that per-acre yield of the crop will be increased, which addresses food supply-chain problems and increases the per-planted-acre profit margin for the farmer. Finally, automated robotic systems will address labor shortages that can limit the size of plantings, the variety of plantings (particularly fruits and vegetables), and the now common occurrence of unharvested crops.

In this paper we present an efficient and robust approach for plant localization. Our method takes as input a 3D point cloud and predicts, accurately, the location of plants along with their height. This type of approach is unique in terms of its simplicity, performance, and extensibility. The method can output plant characteristics that can be found useful for both agricultural monitoring, and in motion planning for a robotic precision agriculture treatment. In Fig. 1 a complete example of the input data along with the predictions of our method can be seen. We chose corn to be the first crop for this application due to its importance both as a food and a feed crop, its non-uniformity of leaf structures, and its heavy use of agrochemicals (in particular, fertilizers with high nitrogen and phosphorous content [4]). The determination of the growth-stage of corn is key information to drive changes in irrigation and the application of agrochemicals [5].

This work's novelty is the 3D based plant localization algorithm that can be used to measure phenotypic traits. Our method is not dependent on training data and is verified in real-world conditions. While other works have done parts of this before, we are one of the first to automatically incorporate row information to make the process fully automatic.

979-8-3503-2365-8/23/$31.00 ©2023 IEEE

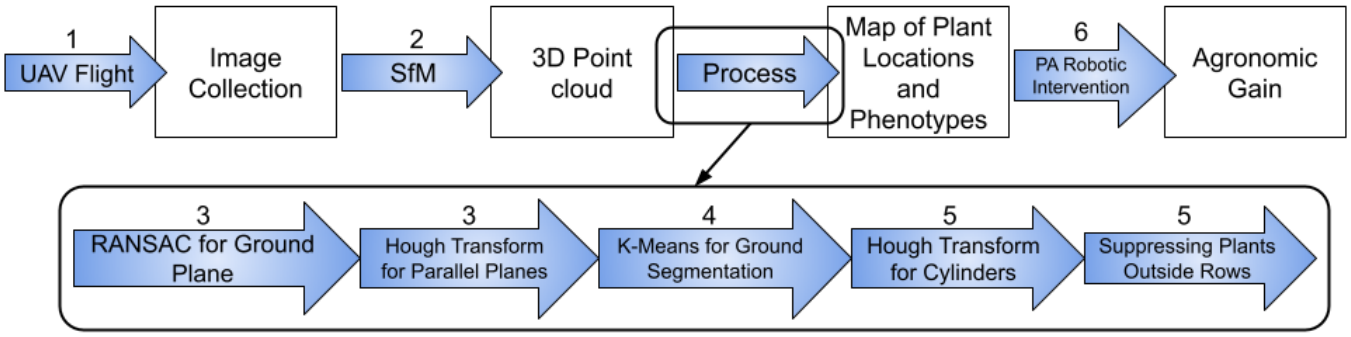


**Fig. 2:** A typical robotic precision agriculture (PA) workflow where our method can be applied is shown in the top flowchart. The bottom flowchart shows the method proposed in this paper as part of the larger workflow.

## II. Related Work

The need to develop automated and large-scale crop monitoring and phenotyping systems [6] [7] has grown in recent years, yet little progress has been made on this front. Advances in the areas of agricultural robotics, remote-sensing, and internet-of-things technologies have increased both the opportunity, and potential rewards of having individual plant information at a large scale [1]; but there is still a lack of automated systems robust to real world conditions.

Aerial surveys of cropland are already widely performed [8]. Works such as [9]–[14] use images from UAVs, and [9]–[12] use UAV images to measure phenotypic traits. While methods like [9] [10] make counts of the plants in the field using image information, the work in [12] attempts to provide plot-level phenotypic measurements rather than individual plant measurements. Convolutional neural networks have also been used to localize and count plants in images from UAVs [10] and ground robots [15], [16], but cannot provide 3D structural phenotype information and requires extensive verified data to train the neural network.

Alternatively, approaches such as [17]–[22], use 3D information rather than images alone. While [17]–[19] use costly LIDAR systems, the research reported in [20], [21], and [22] chooses a 3D reconstruction approach similar to ours. However, [21] only operates in laboratory conditions and does not operate on plants *in situ* and [22] uses a ground-based platform for data collection that will have difficulty navigating a field as crop heights increase. The most similar work to ours is [20] as it operates on UAV collected data of plants in a field setting, and localizes plants and extracts phenotype measurements. While it does incorporate row information, it only uses it to define a known scale in the point cloud and does so through manual input from the user selecting points in two adjacent rows, whereas our method performs these calculations automatically. It achieves a counting accuracy of 88.1% but does so using a sampling based method that is prone to failure when noise is present in the data. Our method achieves an accuracy that is experimentally verified at over 93%.

## III. Method

The process to determine individual plant locations given a dense 3D point cloud has multiple distinct steps. These steps are similar to a traditional data mining pipeline:

1) Collect raw data
2) Extract target data
3) Pre-process the data through filtering and normalization
4) Transform the data through statistical methods
5) Identify patterns in the transformed data
6) Interpret the patterns to gain knowledge

Our research in this paper focuses on **Step 3** through **Step 5**, and Fig 2 demonstrates a full precision agriculture (PA) pipeline along with our focus on the 3D point cloud processing procedure. **Step 1** is a related but separate project that incorporates flight planning and view selection for guaranteeing dense coverage of all plot areas in the point cloud [23]. The analysis of **Step 6** is also a related but separate work that is still under investigation.

**Step 1:** The video is collected from a camera drone that overflies a field plot of row crops, in this initial work, corn plants. The drone flies in a circular path focusing its camera on the center of the corn plot such that only plants unaffected by the prop wash of the UAV are captured in the video. While the circular pattern is not necessarily optimal [23], it provides sufficient density of points after **Step 2**.

**Step 2**: From the drone video, we post-process frames to extract 3D points via structure-from-motion (SfM), which is done using the COLMAP [24] [25] software. At this step, the point cloud is dimensionless.

**Step 3**: The generated point cloud is then analyzed to determine the location of the crop rows, the unit-scale of the dimensions, and the affine transforms required to align the data to the coordinate frame. The location of the crop rows is accomplished using RANSAC [26] to eliminate the majority of the points from the ground plane, leaving only those points that correspond to the crop. We then utilize a Hough Transform to find evenly-spaced parallel planes perpendicular to the ground plane representing each row of the crop. The Hough space for this operation can be seen in Fig. 3. We also rotate the cloud to align (approximately) with a coordinate axis. Alignment is of benefit to the techniques and visualizations used in the following steps. Without alignment, the matrices used in the subsequent storage and processing are sparse. A list data structure would eliminate

the sparsity, but would reduce greatly the overall efficiency due to the non-random access of list structures. Alignment is only $\Theta(N)$ and significantly reduces the impact of $N$.

**Step 4:** Since we have non-data points (matrix padding) and 'dirt' points we need to eliminate them before the voting process. In **Step 3** we identified the ground plane which corresponds to the dirt points; however, there remain some points that were missed by RANSAC due to elevation changes in the field. The ground plane may be sloped with-respect-to the GPS elevation of the drone flight. The tilling of the field also produces variations in the elevation across the field plot. We switch to a K-means clustering approach to determine the rest of the ground plane points by clustering points based on their height above the ground plane found by RANSAC. Once identified, these points do not contribute to the voting scheme we use to determine plant locations.

**Step 5:** The points that remain after **Step 4** are the ones that generally correspond to stalk and leaf points. They will vote for all possible stalk locations that could generate each point. Locations where no stalk exists will tally only leaf information and will have lower vote totals due to the larger number of points generated by stalks.

The stalk is modeled as a vertical or nearly vertical cylinder, the diameter of which is determined by the plant species and variety. These characteristics are treated as hyperparameters and their values are dependent on the application and the type of crop. By using a cylinder diameter at the upper end of the feasible range for a given crop, we can account for cases when the plant stems are not perfectly vertical.

**Step 6:** The analysis of the individual plants and their associated data points will provide estimates of growth stage and potentially the diagnosis of plant deficiencies, plant infections, and plant infestations. This will in turn guide further plant-level diagnosis, agrochemical application, irrigation policies, and yield prediction. In this paper, the height of the plants is measured but other phenotypic traits requiring access to 3D structural information could also be determined.

We present an estimate of individual plant height, but other traits as well as diagnosis/mitigation have yet to be investigated. The effectiveness of our method is validated on two different point clouds.

## IV. Experimental Evaluation

We create a two different dense point clouds derived from drone videos named Cloud 1 and Cloud 2. For brevity, we only detail the evaluation process for Cloud 1 but show results from both point clouds. Fig. 1 shows a visualization of Cloud 1 where the colors of each 3D point are taken from the original video. The points are dimensionless and have ordinal ranges, rounded to thousandths: for X: -57.615 to 60.928; for Y: -60.336 to 56.561; and for Z: -5.086 to 17.144.

We use the RANSAC and the Hough methods to preprocess the target data. The Hough Transform gives parameterized equations for each row. Given *a priori* knowledge of the distance between rows from the planter used on the field, a

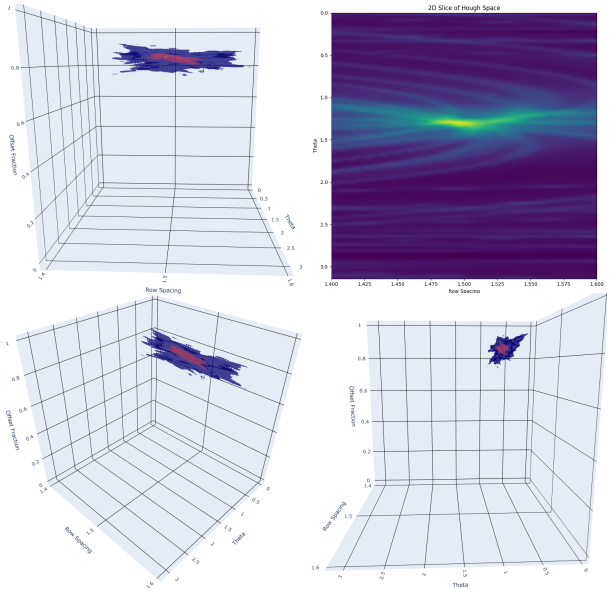


**Fig. 3:** The Hough space showing how points cluster around a given angle, spacing, and offset. Counterclockwise from the upper left, the first three are views of an isosurface plot of the Hough space and the last is a 2D slice shown as an image.

scaling factor can be calculated to provide dimensional units for the points in the data set. This scaling factor is needed to estimate plant stem diameter and overall plant diameter. The row equations will also be used in later steps to cross-check the plant locations resulting from the voting process. This finds strong evidence for five crop rows in the data.

The derived Hough Transform coordinates are still in dimensionless Point Cloud Units (PCUs). Using the 76.2 centimeters (cm) (30-inch) row spacing and adjacent row equations from the Hough Transform shown in Fig. 3, we estimate the scaling factor in cm per PCU as 4.306.

We also align data in this step to reduce the sparse points introduced through padding of the matrix. The actual points only occupy approximately 52% of the X-Y matrix cells before alignment. After alignment, they occupy 97% of the overall X-Y cells. This alignment makes sure both the clustering and the voting will have significantly lower runtime (about twice as fast) since the resulting matrix is more compact.

The system automatically determines the Z-rotation and the X-, Y-, and Z-translations needed for alignment. For this point cloud, the rotation is calculated at 0.806 radians. After the rotation aligns to the desired coordinate axes, a translation is applied to move all points into the first X-Y quadrant. The new ordinal ranges, rounded to thousandths, are for X: 0 to 78.220; for Y: 0 to 91.141; and for Z: 0 to 22.230. As a sanity check on our approach's accuracy, this gives the crop-plot size of 3.368 by 3.924 meters (m) with a maximum height of 95.727 cm.

Estimates for stalk diameter and seed spacing can now be calculated. Stalk diameter is based upon the morphology of the plant species present in the field plot. For the corn variety examined here, stalk diameters are typically 2 to 4.5 cm. Seed spacing is based upon the seed planter's notched

wheels. Both of these values are easily determined by the farmer and essential for our voting scheme and the resulting estimates of plant locations. In PCUs, a typical stem diameter is estimated to be 0.464 PCU to 1.045 PCU. The height is not the actual maximum corn stalk height due to a "hole" in the plot field. This might be an actual hole, or it could be a region of mismatched points from the reconstruction. In either case, once clustered, we can compensate for this hole's depth and the actual plant heights can then be estimated. Height is one of the characteristics that can be used to gain some information of the plant's health and growth stage [4].

Since the points that correspond to the plant structures (stems, leaves, possible ears, tassels, etc.) should be the only points that vote for stem locations, we will remove the points that correspond to the ground plane (dirt points). We use a K-means clustering approach with a total of nine clusters, as we can see in Table I. The number of clusters was determined for Cloud 1 but was also verified to produce good results on Cloud 2 without modification. This number of clusters slightly over-fits the data, but it makes automated separation of the plant and ground points clearer. We cluster on the Z-values in the aligned point cloud. The hole in the field falls into a single cluster, while the rest of the hole and the ground plane points collect in two clusters with very similar means. The clusters give an easily identifiable separation between dirt and plant points due to the magnitude of the clusters. The clustered point cloud can be seen in Fig 4.

After the preceding steps of collection, extraction, processing, and transformation of the point-cloud data, the system can now determine probable plant locations. We can estimate the maximum number of plants using the number of rows, the results of the Hough transforms on those rows, and the plot dimensions. The row lengths are estimated at 341 cm producing a maximum of 19 plants given the known seed drill spacing of 18 cm on the planter. The maximum number of plants would then be about 95 plants, but this is a significant over-estimate due to seed drill misses, seeds that did not germinate, and plot alignment. This estimate assumes that the rows have plants starting at the edge of the plot point-cloud. Probabilistically, this would not be the case and we would expect this to reduce the number of potential planting locations by about 17 (a reduction of 2 locations per row). This gives us a maximum of 85 plants in the plot without compensating for seed drill misses. As we show later, the first row (the top row in the aligned field plot) has only 12 plants, giving 5 likely seed drill misses or plants that did not emerge. This number of misses is relatively high and Fig. 1 shows that this is an extreme outlier in this field plot. Manual verification of the number of plants in the field plot is 87 with 10 of those classified as double seedings, giving a probable 8 seed drill misses or plants that did not emerge (predicted compared to hand-count minus doubles).

**TABLE I:** Results of the K-Means clustering on 3D-point heights.

| Cluster | Mean Height (cm) | # of Points |
|---|---|---|
| 1 | 6.61 | 38 |
| 2 | 17.66 | 775,324 |
| 3 | 20.03 | 1,650,096 |
| 4 | 32.23 | 118,817 |
| 5 | 43.89 | 173,900 |
| 6 | 54.16 | 190,814 |
| 7 | 63.28 | 217,495 |
| 8 | 71.94 | 141,422 |
| 9 | 80.71 | 55,365 |

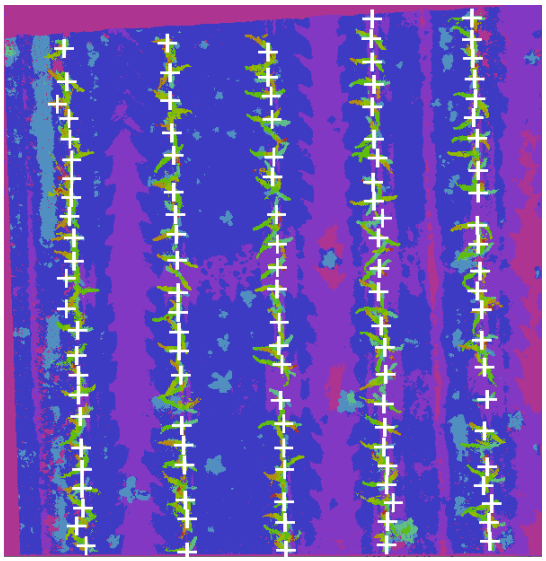

**Fig. 4:** The voting-process reported plant locations annotated on top of the height-clustered points. This image shows Point Cloud 2.

We use a voting scheme based upon vertical point density as shown in Fig. 5. The matrix has the same dimensions as the voting array. The 3D points are projected into the cells of the matrix. The voting array is a $600 \times 700$ matrix where each cell is a possible plant location. The locations are evenly tiled across the entire field plot. Given the prior calculations for the size of the plot, this gives the voting matrix cells a resolution of approximately 0.561 cm. Using the 4.5 cm maximum stalk diameter, a 3D point projected into the X-Y plane casts votes in 69 cells in the voting matrix (a circular mask in a $9 \times 9$ region centered on the 3D point's location in the XY plane).

The resulting values in the voting matrix can be visualized as a gray-scale image. Fig. 5 shows all votes, but since we will suppress votes from points that are clustered as dirt (eliminating around 2.4 of the 3.3 million points), the number of potential locations is greatly reduced from the 420,000 cells in the voting matrix.

Once all the votes have been cast, the locations with the highest votes are processed. For context, in this point cloud the highest voted location is at cell (284, 80) with a total of 7,787 votes. We use an iterative process based upon vote rank.

While there are still cells with votes we perform the following steps:

1) Extract the cell location with the highest vote total,
2) Check that the location falls within 5 cm of one of the row equations identified by the Hough transform in **Step 3**,
3) If the location satisfies a row equation, report the lo-

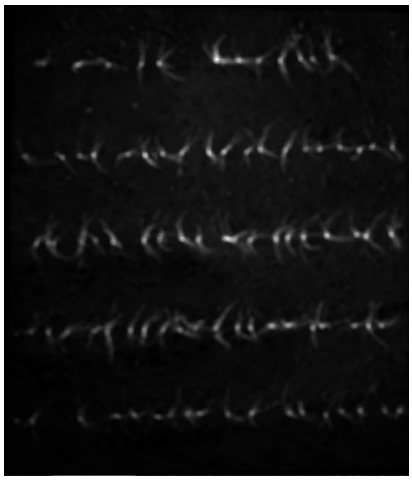

**Fig. 5:** The cylinder votes visualized as an image aligned to the field plot.

cation and zero out the location and the neighborhood cells (described below), and

4) If no row equations are satisfied, zero out the location only.

When a voting-identified location is processed, the algorithm suppresses other nearby high-ranking locations akin to non-maximal suppression. We know that the seed drill spacing is approximately 18 cm, so the suppression does not allow reporting of any location within one half (9 cm) of the seed spacing. This suppression serves multiple purposes. First, a single plant's stalks and leaves produce a small cluster of high-vote cells. We only need one of these locations to be reported. Second, two adjacent plants may have significant leaf overlap which can gain enough votes to report out as a weaker plant location. Indeed, in Table II, a weak plant in the corner of the plot did not gather enough points to be reported, but there were other locations with higher votes that were the result of leaf overlap. Third, double seedings can occur and these may lead to two plants in extremely close proximity to each other. Given the desire to assess plant health for precision treatments, we propose that automated systems need only a single location to be reported out for both plants in the double seeding. The location results of this process are illustrated for Cloud 1 in Fig. 1 and for Cloud 2 in Fig. 4.

To test our methodology, we determined the stalk locations in X and Y (by hand) for the plants in the top row of the plot. This row has several seed drill misses, a seed drill double (near the right end of the row), and some inconsistent seed spacing. We take the point cloud units of these ground-truth locations and follow the previously calculated affine transformations to align and to scale these ground truth locations. This calculates the true locations in the coordinate system of the voting matrix. A total of twelve plants appear in this row; however, one of these plants is considered a double seeding. This plant is not identified by the voting system.

Table III shows the actual locations determined by hand and the voting array locations that were identified by our algorithm. We also calculate the positional errors and convert them to cm. Since a single voting matrix cell is $1.46 \times 1.46$ cm, even a single cell location error in each axis will result in up to a 2.06 cm error.

To characterize further the approach's performance, we provide in Table II a breakdown of each row's plants and their detections. We expect double seedings to only count as one plant. Overall, there are 87 actual plants with 10 instances of double seedings (plants closer than one-half of the seed drill's spacing). This gives 77 expected plant detections. The system reported 72 of these plants for a success rate of 93.5%.

Finally, we can estimate the mean plant height from the last cluster minus the ground plane (at approximately 18.10 cm). This gives us average plant heights of approximately 63 cm and a maximum plant height (based upon the prior maximum Z) of about 78 cm. Given the reported locations and a simple maximization of the Z values in the neighborhood of each of those locations, finding the height of each identified plant is trivial. This is one of the critical measures for determining corn growth stage [4]. Fig. 6 shows the distribution of plant heights found in the two point clouds processed with our method. The heights can also be seen in Fig. 1 where their accuracy can also be qualitatively evaluated as we do not have ground truth height measurements to compare with.

Satisfied with the accuracy of the data from the method, we ran the system on another corn plot point cloud (Cloud 2) that contains 4,343,165 3D points. This plot also contains 5 rows but is slightly larger than Cloud 1. This particular plot is a good challenge for our methodology as the plants are smaller and will therefore produce fewer votes both during row identification and during plant localization. The till of the dirt is less even, there are non-corn plants present in the data (weeds), and the number of plants is higher, as is the number of data points.

The method calculates a plot size of 4.067 by 4.166 m. The till of the dirt has significant rounding along the plot rows and shows seed drill patterns that are observable in the clustering (see Fig. 4). Again, Cluster 3 is identified as the ground plane height such that the estimated maximum plant height is 37.4 cm. Indeed, the point-cloud data definitely shows plants that are smaller and at an earlier growth stage than in Cloud 1. The calculated lengths of the identified rows and the 18 cm seed drill spacing produces approximately 25 plants per row

**TABLE II:** The hand-counted plants per row (numbered from top to bottom) minus the hand-counted double seedings gives an expected number of detections. The algorithm detection shows accuracy/miscounts.

| Row | Hand Counted Plants | Hand Identified Doubles | Expected Number of Detections | Number Actually Detected |
|---|---|---|---|---|
| 1 | 12 | 1 | 11 | 11 |
| 2 | 18 | 1 | 17 | 16 |
| 3 | 19 | 3 | 16 | 14 |
| 4 | 20 | 2 | 18 | 17 |
| 5 | 18 | 3 | 15 | 14 |
| Totals | 87 | 10 | 77 | 72 |

**TABLE III:** Prediction error of our voting method per plant.

| Plant | Actual Position (X,Y) | Voted Position (X,Y) | Error (cm) |
|---|---|---|---|
| 1 | (60.432, 84.624) | (58, 86) | 4.09 |
| 2 | (109.471, 89.420) | (110, 88) | 2.22 |
| 3 | (154.154, 88.644) | (155, 88) | 1.55 |
| 4 | (202.634, 91.405) | (204, 87) | 6.74 |
| 5 | (228.981, 88.728) | (232, 86) | 5.95 |
| 6 | (312.835, 83.086) | (316, 80) | 6.46 |
| 7 | (348.403, 79.349) | (350, 82) | 4.53 |
| 8 | (374.585, 83.301) | (379, 82) | 6.73 |
| 9 | (409.092, 81.396) | (412, 83) | 4.86 |
| 10 | (438.892, 82.612) | (441, 79) | 6.12 |
| 11 | (486.669, 76.613) | (486, 77) | 1.13 |
| Average | - | - | 4.58 |

for a maximum of 125 possible plants (discounting double seedings).

The presence of weeds in the point cloud clearly affects the nearby corn plant identification. This is observable in both the second and the last rows from the left in Fig. 4. Fig. 6 gives the histogram of plant heights found in both point clouds.

Overall, the multi-step method we have created is efficient. At each step, the asymptotic runtime for the underlying algorithm stays low. In **Step 2** COLMAP takes quadratic time based on the number of views (exhaustive matching) but this process is easily parallelizable and runs quickly on GPUs [24], [25]. The point clouds used in this work used roughly 100 different camera views causing the reconstruction to be the most expensive part of the method in terms of run-time by several orders of magnitude, averaging about 4 hours each. The RANSAC and Hough algorithms used in **Step 3** take constant and linear time respectively.

The K-means clustering (**Step 4**) runs in linear time since we limit the number of iterations to a constant (a maximum of 100). This limit is never reached even though we require convergence to a small entropy term (0.1 average cluster entropy from the *n-1*st to the *n*th iteration). In fact, given our selection of 9 height clusters that gives a slight overfit, the observed iterations for clustering was 4 on Cloud 1 and 3 on Cloud 2. We lowered the entropy threshold by an order of magnitude (to 0.01) to test this claim and the clustering converges in 11 iterations on Cloud 1, with most of the point migration occurring from Cluster 3 to Cluster 2. These are both clusters associated with the ground plane (e.g., dirt).

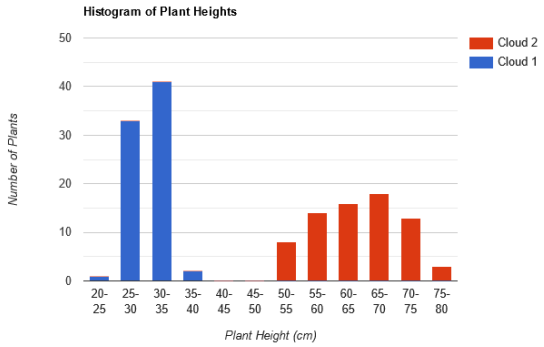


**Fig. 6:** Histogram of plant heights detected in each point cloud.

## V. Conclusion

We have presented a multi-stage algorithm that accurately reports corn plant locations in a point cloud that was generated by a simple camera drone overflight. The simplicity of the raw data collection and the fully autonomous generation of individual plant data makes this system both inexpensive and easy to use. These are essential traits for any system to be adopted broadly by farmers. Complex systems that require skilled human interaction are simply not viable at scale.

The *a priori* data required by the system are measurements that farmers already know from their planting equipment: the row spacing of the planter; the seed spacing of the notched disks on the seed drill; and the average size of the crop stalk. These values vary by farmer, by crop, and by equipment, but they are always known and available. All other measurements and constants are derived autonomously by our method.

We have verified the approach through testing on actual field plot data and hand-calculated ground truth. In this testing, the system accurately reported plant locations within a few cm of the actual position. The mean error was 4.58 cm. It also located 93.5% of the potential plant locations (discounting double seedings). We confirmed these initial results using a second, larger point cloud (Cloud 2) that presents several challenges to the proposed method. Our approach is still effective on this more challenging set of data.

## VI. Acknowledgements

This work is supported by the Corn Growers Association of MN, the Minnesota Robotics Institute (MnRI), Honeywell, and the National Science Foundation through grants #CNS-1439728, #CNS-1531330, and #CNS-1939033. USDA/NIFA has also supported this work through the grant 2020-67021-30755.

## References

[1] N. Zhang, M. Wang, and N. Wang, "Precision agriculture—a worldwide overview," *Computers and Electronics in Agriculture*, vol. 36, no. 2, pp. 113–132, 2002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169902000960

[2] B. T. Nolan, K. J. Hitt, and B. C. Ruddy, "Probability of nitrate contamination of recently recharged ground waters in the conterminous united states," *Environmental Science and Technology*, vol. 36, no. 10, pp. 2138–2145, 1981.

[3] S. F. Johnson, "Methemoglobinemia: Infants at risk," *Current Problems in Pediatric and Adolescent Health Care*, vol. 49, no. 3, pp. 57–67, 2019.

[4] L. Abendroth, R. Elmore, R. G. Hartzler, C. McGrath, D. S. Mueller, G. P. Munkvold, R. Pope, M. E. Rice, A. E. Robertson, J. Sawyer, *et al.*, "Corn field guide," *Ames (IA): Iowa State Univ Extension and Outreach Publications*, 2009.

[5] J. J. Hanway *et al.*, "How a corn plant develops." *Special Report, 38*, 1966.

[6] R. T. Furbank and M. Tester, "Phenomics – technologies to relieve the phenotyping bottleneck," *Trends in Plant Science*, vol. 16, no. 12, pp. 635–644, 2011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1360138511002093

[7] 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.

[8] Y. Shi, J. A. Thomasson, S. C. Murray, N. A. Pugh, W. L. Rooney, S. Shafian, N. Rajan, G. Rouze, C. L. Morgan, H. L. Neely, *et al.*, "Unmanned aerial vehicles for high-throughput phenotyping and agronomic research," *PloS one*, vol. 11, no. 7, p. e0159781, 2016.

[9] F. Gnädinger and U. Schmidhalter, "Digital counts of maize plants by unmanned aerial vehicles (uavs)," *Remote Sensing*, vol. 9, no. 6, p. 544, 2017.

[10] A. Karami, M. Crawford, and E. J. Delp, "Automatic plant counting and location based on a few-shot learning technique," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 5872–5886, 2020.

[11] N. Chebrolu, T. Läbe, and C. Stachniss, "Robust long-term registration of uav images of crop fields for precision agriculture," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3097–3104, 2018.

[12] R. Makanza, M. Zaman-Allah, J. E. Cairns, C. Magorokosho, A. Tarekegne, M. Olsen, and B. M. Prasanna, "High-throughput phenotyping of canopy cover and senescence in maize field trials using aerial digital canopy imaging," *Remote Sensing*, vol. 10, no. 2, 2018. [Online]. Available: https://www.mdpi.com/2072-4292/10/2/330

[13] C. Nived, L. Philipp, L. Thomas, and S. Behnke, "Robot localization based on aerial images for precision agriculture tasks in crop fields," in *Robotics and Automation (ICRA), 2019 IEEE International Conference on*, 2019.

[14] N. Chebrolu, T. Läbe, and C. Stachniss, "Robust long-term registration of uav images of crop fields for precision agriculture," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3097–3104, 2018.

[15] P. Lottes, J. Behley, N. Chebrolu, A. Milioto, and C. Stachniss, "Joint stem detection and crop-weed classification for plant-specific treatment in precision farming," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 8233–8238.

[16] 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.

[17] G. Alenya, B. Dellen, and C. Torras, "3d modelling of leaves from color and tof data for robotized plant measuring," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 3408–3414.

[18] S. Paulus, H. Schumann, H. Kuhlmann, and J. Léon, "High-precision laser scanning system for capturing 3d plant architecture and analysing growth of cereal plants," *Biosystems Engineering*, vol. 121, pp. 1–11, 2014.

[19] Y. Shi, N. Wang, R. K. Taylor, W. R. Raun, and J. A. Hardin, "Automatic corn plant location and spacing measurement using laser line-scan technique," *Precision Agriculture*, vol. 14, no. 5, pp. 478–494, 2013.

[20] 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.

[21] A. Paproki, X. Sirault, S. Berry, R. Furbank, and J. Fripp, "A novel mesh processing based technique for 3d plant analysis," *BMC plant biology*, vol. 12, no. 1, pp. 1–13, 2012.

[22] 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.

[23] A. Bacharis, H. J. Nelson, and N. Papanikolopoulos, "View planning using discrete optimization for 3d reconstruction of row crops," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 9195–9201.

[24] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[25] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, "Pixelwise view selection for unstructured multi-view stereo," in *European Conference on Computer Vision (ECCV)*, 2016.

[26] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 2002.