Project Title: Decision Support Tool for Precision Orchard Management

Report Type: Final Project Report

Primary PI: Joseph Davidson
Organization: Oregon State University
Telephone: 541-737-9193
Email: joseph.davidson@oregonstate.edu
Address: 204 Rogers Hall
Address 2: 200 SW Monroe Avenue
City/State/Zip: Corvallis, OR 97331

Co-PI 2: Cindy Grimm
Organization: Oregon State University
Telephone: 541-737-4914
Email: grimmc@oregonstate.edu
Address: 204 Rogers Hall
Address 2: 200 SW Monroe Avenue
City/State/Zip: Corvallis, OR 97331

CO-PI 3: Ashley Thompson
Organization: Oregon State University
Telephone: 541-296-5494
Email: ashley.thompson@oregonstate.edu
Address: 400 E Scenic Dr Suite 2.27
Address 2:
City/State/Zip: The Dalles, OR 97058

Co-PI 4: Manoj Karkee
Organization: Washington State University
Telephone: 509-786-9208
Email: manoj.karkee@wsu.edu
Address: WSU Prosser – IAREC
Address 2: 24106 N. Bunn Road
City/State/Zip: Prosser, WA 99350

Cooperators: Dave Allan (Allan Brothers Fruit Co.)

Project Duration: 3 Year

Total Project Request for Year 1 Funding: $73,569
Total Project Request for Year 2 Funding: $77,335
Total Project Request for Year 3 Funding: $71,596

Other related/associated funding sources: None

WTFRC Collaborative Costs: None

Budget 1
Primary PI: Joseph Davidson
Organization Name: Oregon State University/Agricultural Research Foundation
Contract Administrator: Charlene Wilkinson
Telephone: (541) 737-3228
Contract administrator email address: charlene.wilkinson@oregonstate.edu

Item	**2020**	**2021**	**2022**
Salaries	$31,331.00	$32,271.00	$26,622.00
Benefits	$8,311.00	$9,206.00	$8,162.00
Wages			
Benefits			
RCA Room Rental			
Shipping			
Supplies	$2,986.00	$4,000.00	$4,000.00
Travel	$3,000.00	$3,000.00	$3,000.00
Plot Fees			
Miscellaneous			
Total	$45,628.00	$48,477.00	$41,784.00

[1]Salaries include a Graduate Research Assistant on a 12-month, 0.49 FTE appointment in years 1 and 2, and a 9-month, 0.49 FTE appointment in year 3. Salaries also include 0.25 months per year for Joe Davidson and Cindy Grimm.
[2]Leaf samples are included in the supply budget.
[3]Travel budget is requested to support mileage and lodging for data collection and field experiments.

Budget 2
Co PI 2: Manoj Karkee
Organization Name: Washington State University
Contract Administrator: Katy Roberts
Telephone: 509-335-4564
Contract administrator email address: katy.roberts@wsu.edu

Item	**2020**	**2021**	**2022**
Salaries	$17,840.00	$18,554.00	$19,296.00
Benefits	$5,101.00	$5,304.00	$5,516.00
Wages			
Benefits			
RCA Room Rental			
Shipping			
Supplies	$4,000.00	$4,000.00	$4,000.00
Travel	$1,000.00	$1,000.00	$1,000.00
Plot Fees			
Miscellaneous			
Total	$27,941.00	$28,858.00	$29,812.00

[1]Travel budget is requested to cover the mileage for field experiments.

Introduction: The standard practice of broad-acre orchard management does not result in targeted actions that are optimal for individual trees – this reduces the impact of management decisions and wastes resources while falling short on achieving the yield and quality potential of individual blocks. **Our team's overall goal was to improve fruit quality and yields by managing individual trees through a combination of automated sensing, learning algorithms, decision support tools, and precision application with variable rate technology**. While for this project we focused on matching nitrogen (N) fertilizer to N demand, our long-term vision is to extend this framework for farming at the tree level to other orchard management decisions (e.g. plant growth regulators, root pruning, tree pruning, chemical thinning). The conceptual framework that we developed for precision N application is shown in Fig. 1 and included the following sequence of activities:

1. Build a site map of individual trees (performed once at the beginning of the project)
2. Use non-contact sensing to estimate tree nutrition (performed annually)
3. Recommend tree-specific fertilization plans using decision support tools incorporating machine learning
4. Apply variable rate N using real-time vehicle localization and precision technology
5. Use historical data to improve the performance of the decision support tool

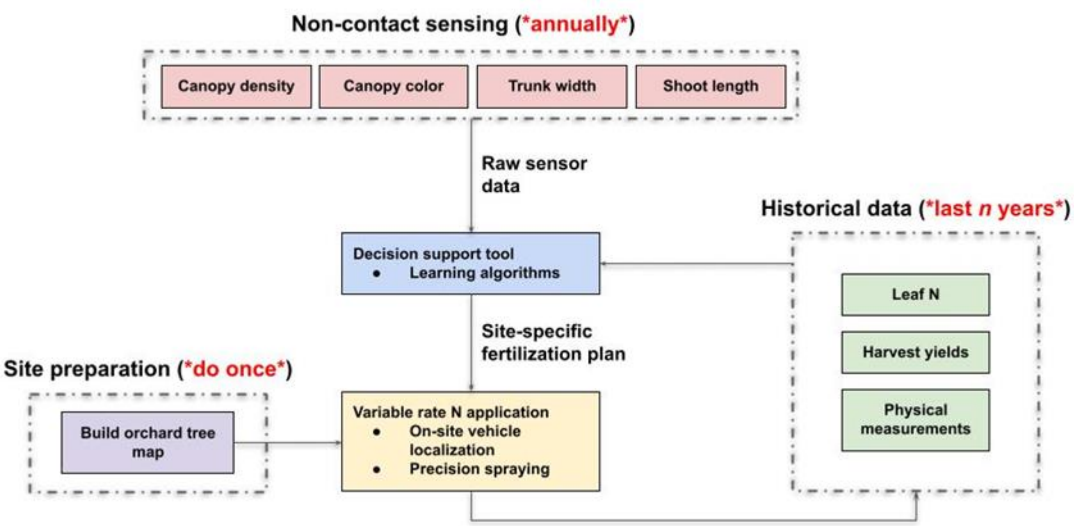

Figure 1. Project framework. A detailed tree map was developed for the site at the beginning of the project. Raw sensor data on various orchard parameters was used as input to a learning algorithm that provides precision fertilization plans. Onsite vehicle localization was used to execute precision application of nitrogen. Historical data on destructive leaf N measurements, horticultural measurements, harvest yields, etc. was used to tune the learning algorithm.

To implement the framework shown in Fig. 1, we created the following 3 specific research objectives:

1. Develop a ground vehicle-mounted sensor system that *i*) maps the geographic location of individual trees within an orchard block; and *ii*) measures plant parameters (e.g. shoot vigor, trunk cross-sectional area, and fall leaf color) to estimate the N status of individual trees
2. Develop a decision support tool that recommends N application levels per tree and tracks the tree's long-term response
3. Develop and demonstrate a proof-of-concept precision spray system that localizes the vehicle with the orchard map, identifies the neighboring trees, and then selectively applies the desired level of N within the root zone

This final report summarizes research results over the performance period of January 2020 – November 2023. *The most significant findings from the project include the following*:

- Trees in the test plot used for the long-term study were more likely to have excess leaf N than be N-deficient.
- Consumer grade RGB-D sensors and state-of-the-art deep learning models can be used for automatic measurement of trunk cross sectional area.
- Normalized canopy area is highly correlated (negatively) to the target N application rate for the upcoming season. The use of normalized canopy area as a measure of canopy vigor and an indicator of the N need of individual trees is promising.
- The yellowness index of a tree at different weeks and the pattern in which they are changing can be a potential indicator of the N status in the tree. Also, weeks 3-4 can be a good time to differentiate between high N and low N trees as the trees with low N start to change color.
- Trunk width can be used to precisely localize a robot within an orchard without the need for GPS.

Objective 1: Orchard mapping & nitrogen sensing

<u>Task 1 – Tree trunk detection (OSU lead, WSU participant)</u>

We created a method that uses an RGB-Depth camera (Intel RealSense D435) to automatically calculate trunk width, which is needed for robot localization and is also an indicator of tree nutrition. The first step to estimate trunk width is to pass the RGB image through an image segmentation model (Deep Neural Network) that determines where the trunk is in the image. The depth image can then be used to determine the distance between the trunk and the camera, which is used to calculate the width of the tree.

Instance segmentation: The segmentation model we adopted was YOLOv8. To train it, 1090 images from the Jazz test block (Yakima Valley Orchards, Prosser, WA) were labeled with post and tree classes. Images from multiple seasons were included to increase model robustness (e.g. dormant season, bloom, and shortly before harvest). The segmentation model performs best when it is used in the test block with the camera positioned similarly as in the training images. Thorough tests were not conducted for other camera positions or in other orchards, but preliminary findings show that its performance notably diminishes. However, training a new segmentation model for images taken from an alternate perspective, or from a different orchard, is relatively simple and very easy to swap out in the width estimation pipeline. Figure 2(left) shows an example of the segmentation. Processing time was 14.6ms per image on our test computer.

Figure 2. (Left) A typical segmentation of an image with a post and trunk. (Right) the width calculation method. The pink lines are the left and right sides of the object, the blue lines are the centerlines, the red lines are the initial width predictions, and the green lines are the corrected widths. The width is shown at three locations for illustration purposes.

Width estimation: Once a trunk or post has been found in the image, we use the aligned depth image to calculate its width. First, the distance to the trunk is found by taking the median depth to the trunk. Next, we calculate an initial estimate of width in pixels by finding the distance between the right and left most pixels of the trunk for each row of the image. However, this estimate is only accurate if the trunk is perfectly vertical, so the width must be adjusted based on the angle of the trunk in the image. To do this, the center line of the trunk is approximated as the point between the rightmost and leftmost point, and then the local angle of the trunk is found by calculating the angle between every 15th pixel along the centerline (Figure 2(right)). We then multiply the width by the sine of the local angle to correct it and use the 40th percentile width as the width estimate. Finally, there is a final correction based on the proximity of the trunk with the edge of the image. We observed that widths for trunks near the edge of an image are overestimated, due to distortions in the camera, so we developed a regression model to correct for this effect. The performance of the final system vs. ground truth human measurements is shown in Figure 3. Mean absolute error was 2.56mm with a standard deviation of 2.39mm. Average processing time to calculate the width was 37.2ms.

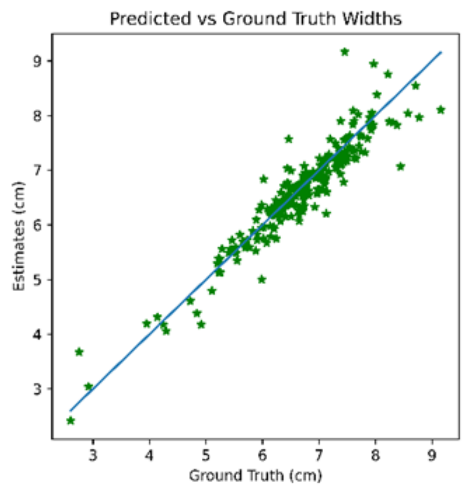

Figure 3. The predicted vs. ground truth measurements for 224 trees using our width measurement system.

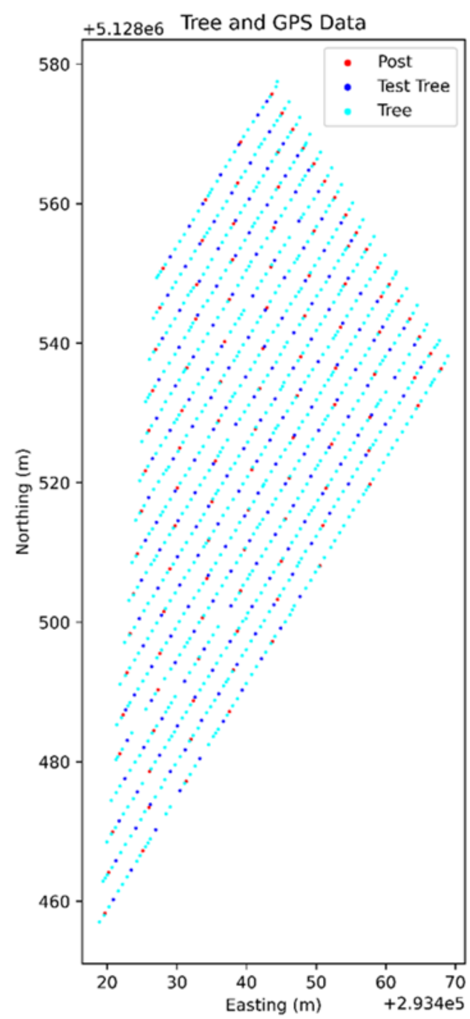

Figure 4. Corrected map of orchard test block.

Task 2 – Orchard mapping (OSU lead, WSU participant)

At the beginning of the project, we created an accurate map of the orchard test block. To create the map, we installed both an RTK GPS (elevated to the top of the canopy on a boom) and an RGB-D camera (mounted on the side) on a utility vehicle and collected data while driving through each row of the site. The data was then postprocessed to create an initial map using our image segmentation pipeline and the GPS recordings. Some corrections had to be made manually, primarily due to either poor image segmentations or errors when matching trunks between successive images from the video stream. Figure 4 shows the final corrected map of the orchard test block. The 200 treatment trees that we used for long term studies are shown within the map.

Task 3 – Nitrogen measurements and non-contact sensing (WSU lead, OSU participant)

We collected leaf samples for the 200 test trees each of the 4 years of the project. Figure 5 shows the results of the mineral analysis. In general, the treatment trees had excess leaf nitrogen as opposed to too little, and the nitrogen status appeared to remain relatively consistent year to year.

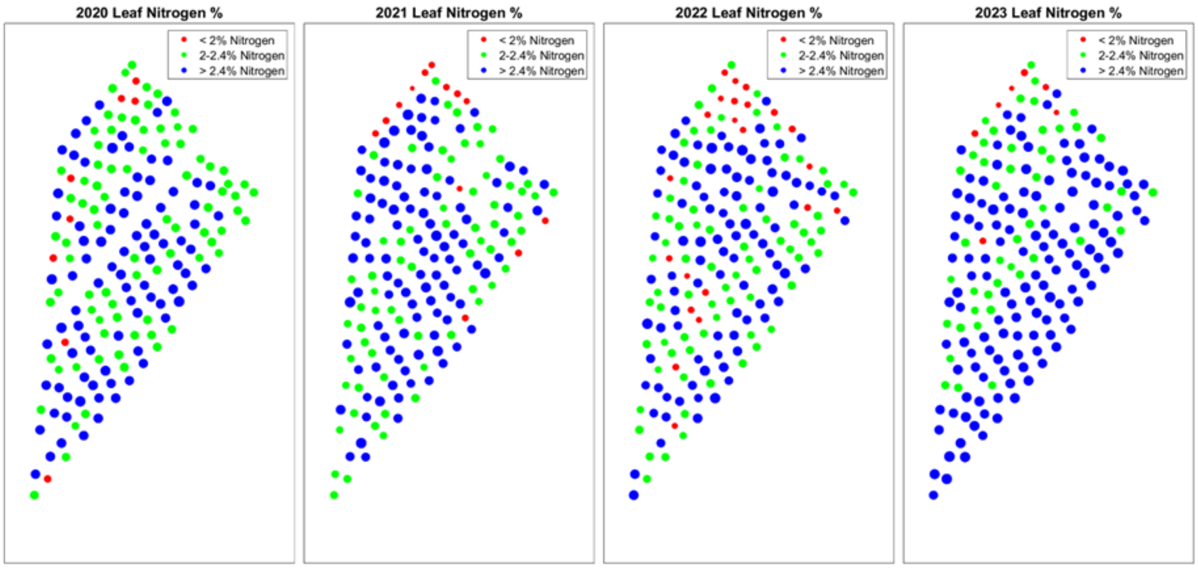

Figure 5. Leaf nitrogen percentage of test trees for 2020-2023. The location of the dot for each treatment tree corresponds to its actual position in the orchard and the size and color correspond to the leaf nitrogen content.

Normalized Canopy Area Estimation

Methods: The dataset for this study was collected during the end of the growing season (late July to early August) in 2021, 2022, and 2023 using a commercial Zed2i sensor mounted on a ground vehicle (Fig. 6). The camera was positioned approximately 7.5 ft above ground and 5.0 ft from the tree trunk so that the entire canopy was visible in individual frames. The vehicle was driven straight in between the tree rows, maintaining an approximately constant distance between the camera and the trees.

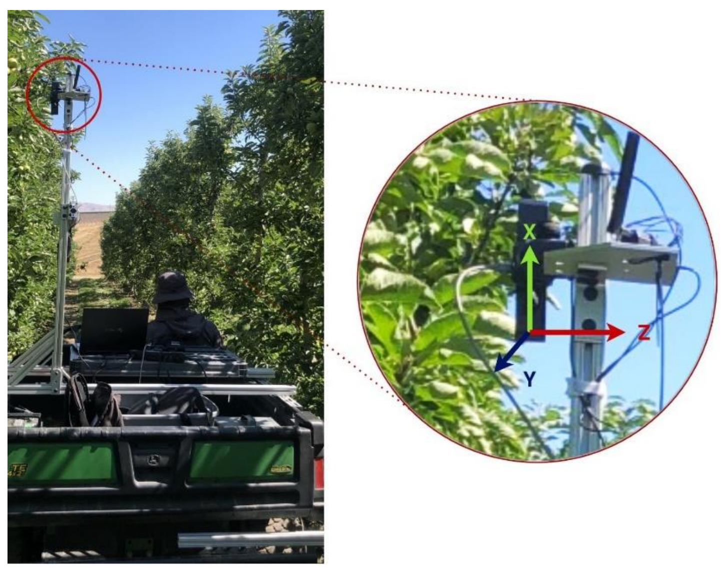

Figure 6. Data collection setup that includes a utility vehicle with Zed2i camera (Stereolabs, France) for image acquisition. The camera and its axes are zoomed in on the right.

After extracting the colored point cloud of the desired frames with the test tree, all points with missing depth information (which could be caused by occlusion or other sensor limitations) were removed followed by color, depth, and height thresholding to remove points from the ground, background, and sky. The point cloud at this step was still noisy, and therefore a radius-based outlier removal technique was used to remove one or a small group of points that were not within a specific distance from the rest of the points in the point cloud. A radius of 20 mm and a total number of points of 50 were used as thresholds to implement this technique using the open3D library, which meant that if there were 50 or fewer points in a spherical neighborhood of 20 mm radius, those were removed from the point cloud. After the radius-based outlier removal, a camera-axis-aligned bounding box was fitted to the point cloud. The bounding box was set to an overall height of 3.5 m, a width of

1.2 m, and a depth of 1.6m. This bounding box ensured that the sample tree was positioned inside the box and that the box size remained the same for all the trees used in this study. The close spacing of trees in the high-density orchard makes it difficult to delineate the boundaries of individual trees. By considering the overall width of 1.2m (4ft) (which is the spacing between trees), we can safely assume that the foliage from the neighboring trees does not significantly affect the calculation of the normalized canopy area of the trees. The segmented point cloud was then projected back to a 2D color image using the indices of the points in the point cloud to obtain the segmented foreground tree image.

To assess the performance of the canopy segmentation technique, the original RGB images for 20 individual tree canopies were manually annotated and compared against the results of the automated segmentation system. The outer boundary of the manually annotated tree canopies was outlined with red, and the holes within the outer boundary of the canopies were outlined with blue (Fig. 7a). The area covered by the outside and inside boundaries was then subtracted to get the total mask (and area in the number of pixels) for the target tree canopies. The canopy areas of individual trees segmented by the proposed algorithm were then compared against the area calculated with the manual annotation (Fig. 7b). The performance of the segmentation algorithm was then evaluated using the Precision, Recall, and F1 matrices as defined by equations 1, 2, and 3, respectively.

$$Precision = True\ Positive \div (True\ Positive + False\ Positive) \quad (1)$$

$$Recall = True\ Positive \div (True\ Positive + False\ Negative) \quad (2)$$

$$F_1 = 2 * Precision * Recall \div (Precision + Recall) \quad (3)$$

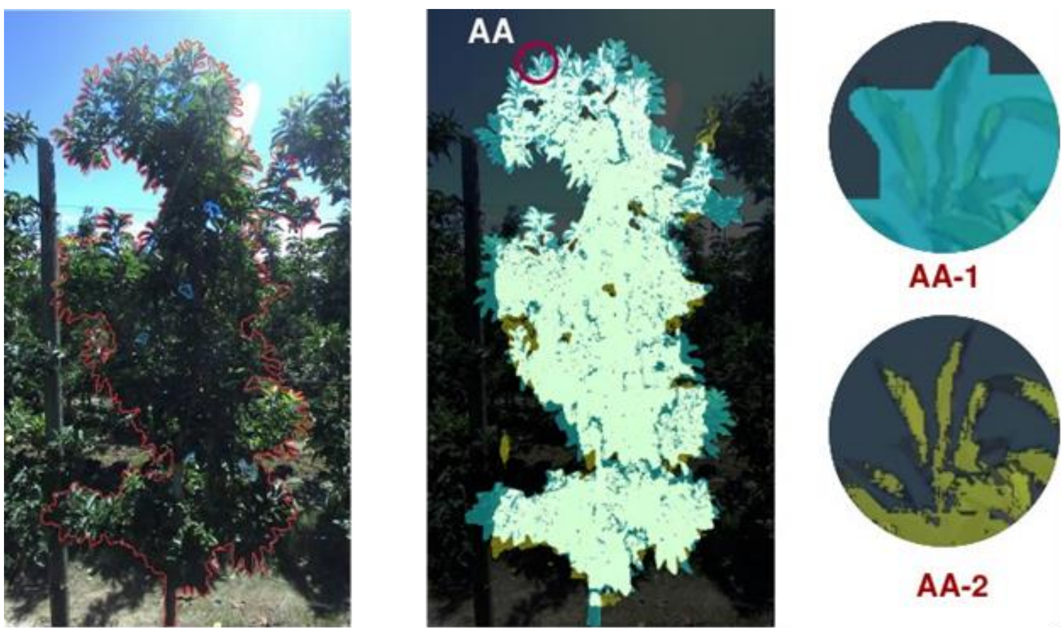

(a) (b)

Figure 7. a) Manually annotated canopy boundary of a sample tree; the red line represented the outside boundary, and the blue lines represented the holes within the overall canopy boundary; b) Segmentation results (yellow) on top of the manually annotated (blue) canopy of the sample tree. White-colored regions represented portions present in both. Sections AA-1 and AA-2 represented the manual annotation and automated segmentation masks for section AA.

The segmented image obtained in the previous step was then used to create a binary mask for the target trees. A rectangular bounding box was fitted outside the binary mask of the segmented tree, which gives the maximum area the canopy could cover. The normalized canopy area was then calculated inside the bounding box using equation 4, which measures how much of the total possible area of the tree is covered by the actual canopy.

$$Normalized\ Canopy\ Area\ = Pixels\ occupied\ by\ vegetation\ \div\ Total\ number\ of\ pixels \quad (4)$$

The normalized canopy area was then correlated with the N application rate recommended by 4 experts based on their assessment of tree vigor. The experts were requested to evaluate the trees and recommend a N application rate on a scale of 0-50 lbs per acre (0 – 56 kg/ha) at an increment of 10 lbs per acre. Two of the experts were growers/orchard managers and were directly involved in orchard management for more than 30 and 15 years, respectively. The other two were from academia and were involved in research, education, and outreach on tree fruit physiology, including nutrition management. The experts independently evaluated the trees after harvest in the same season, as per their availability, before the trees started losing leaves to ensure that the canopy growth from the current season was still completely visible. These experts qualitatively assessed the trees in terms of their vigor (e.g., shoot growth, canopy area) and any other apparent factors (e.g., trunk diameter) they would use in their day-to-day decision-making and incorporated their experience to recommend a N application rate. A higher N application rate means the experts found the trees to have lower vigor than desired for target fruit yield and quality and vice versa. A sub-sample of 55 trees from a total of 199 treatment trees was chosen randomly for the experts' evaluation.

Results & Discussion: A normalized confusion matrix was created to qualitatively assess the performance of the segmentation method (Fig. 8). In this case, false negatives are foreground pixels classified as background, and false positives are background pixels classified as foreground. True positives and true negatives are correctly classified foreground and background pixels. The algorithm achieved a precision of 0.79, recall of 0.77, and an F1 score of 0.78 and was able to segment the foreground tree properly. These results suggest that parts from the ground, background rows, and neighboring trees were properly removed from the analysis with minimal errors.

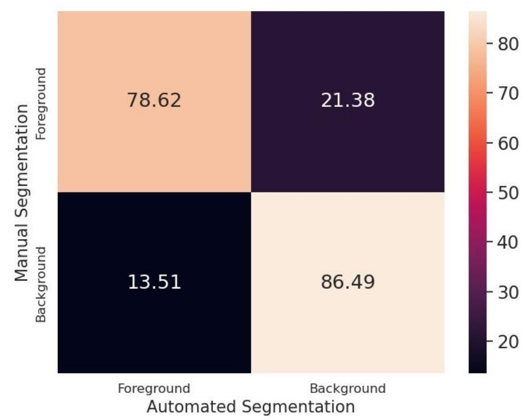

Figure 8. Normalized confusion matrix comparing the automated canopy segmentation with the manually segmented ground truth canopies.

A generally linear relationship existed between the normalized canopy area and the experts' N recommendation (Fig. 9). A correlation coefficient of -0.86, -0.84, -0.96, and -0.78 was obtained between the normalized canopy area and the N application recommendations provided by Experts 1, 2, 3, and 4, respectively. The high negative correlation coefficients among the experts showed that the normalized canopy area is highly correlated (negatively) with the target nitrogen application rate for the upcoming season. The results also indicate that the use of normalized canopy area as a measure of canopy vigor and an indicator of the nitrogen need of individual trees is promising.

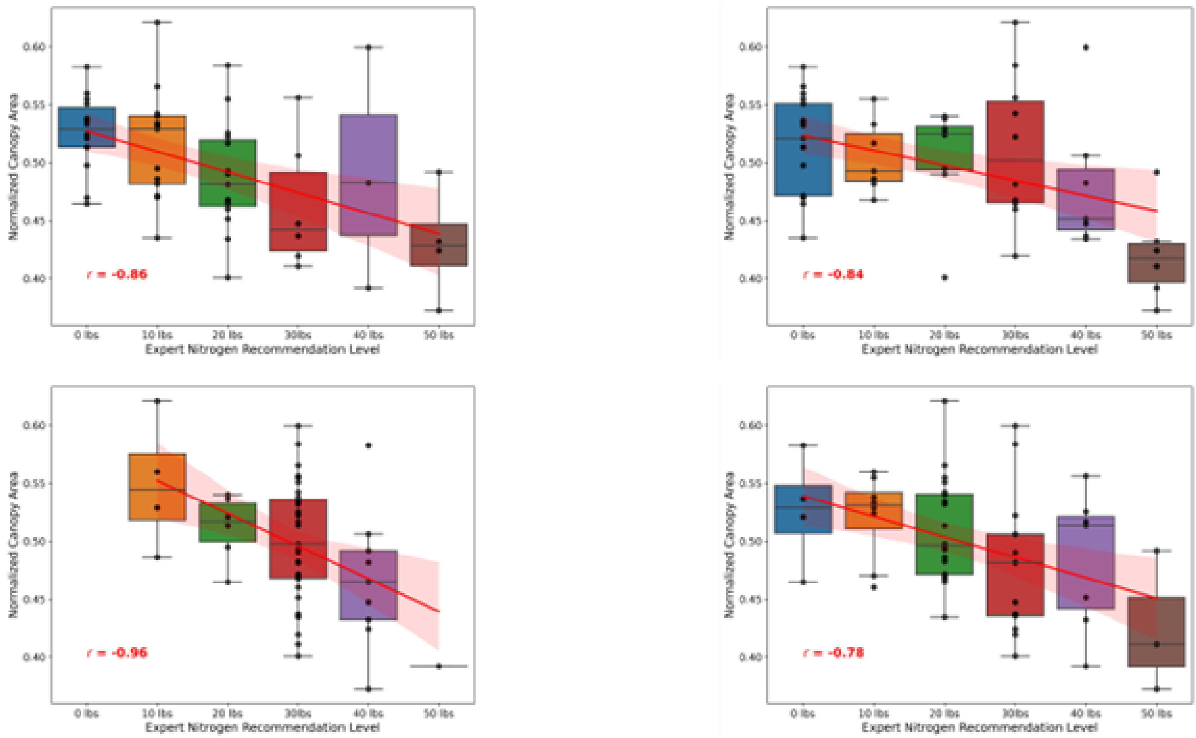

Figure 9: Relationship between normalized canopy area and N application level recommended by (Top left) Expert 1, (Top right) Expert 2, (Bottom left) Expert 3, and (Bottom right) Expert 4. The red line shows the regression plot for the expert nitrogen recommendation level and the median normalized canopy area at each recommendation level for each expert.

<u>**Temporal leaf color assessment**</u>

Methods: Another parameter of interest is the temporal change in leaf color during the fall. For this study, we collected images of the leaves' color change from green to yellow over 6 weeks starting on October 8, 2021 (the leaves froze before showing any color change in 2022, and more data is being collected in 2023). Figure 10 shows the data collection setup and color during the different weeks of the study for one of the sample trees. The point cloud obtained from the camera was thresholded using color and depth thresholds and downsampled uniformly at a 10:1 ratio. The downsampled point cloud was then clustered using a hierarchical clustering technique on the CIE-L*a*b color space. A hierarchical K-means clustering was used to first group the points into 20 clusters. A threshold in both a* and b* spaces for the group centers was applied to merge the classes into 3 final clusters: Yellow, Green, and Trunk. The Yellow cluster included the foliage that had turned yellow, the Green cluster included foliage that was still green, and the Trunk cluster included the remaining points from the trunk, branches, some brown leaves, and soil from the background (and the leaves that had turned red on a few trees). The final output from the clustering algorithm included three clusters: Green cluster (c_g), Yellow cluster (c_y), and Trunk cluster (c_t). Figure 11(left) shows the result of the clustering technique for one of the sample trees where yellow, green, and trunk clusters belong to c_y , c_g , and c_t respectively.

After the grouping of points into 3 clusters/classes, the Yellow (c_y) and Green (c_g) classes were used to calculate the yellowness index of each tree, a metric that we defined to indicate what fraction of the foliage is yellow as compared to green. The yellowness of each tree was calculated using equation 5.

$$Yellownes\ Index\ =\ (y-g) \div (y+g) \qquad (5)$$

where, y = number of pixels/points in the Yellow Cluster, c_y and g = number of pixels/points in the Green Cluster, c_g.

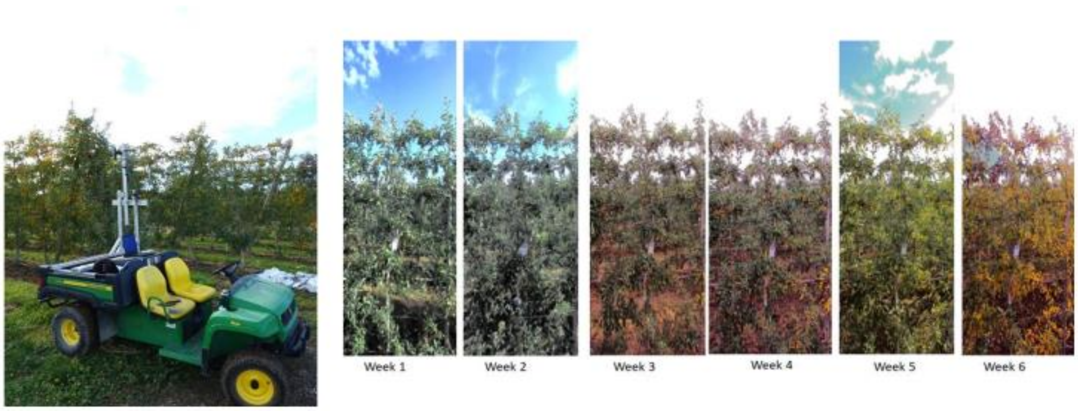

Figure 10: Data collection setup and color change during the six weeks of the study.

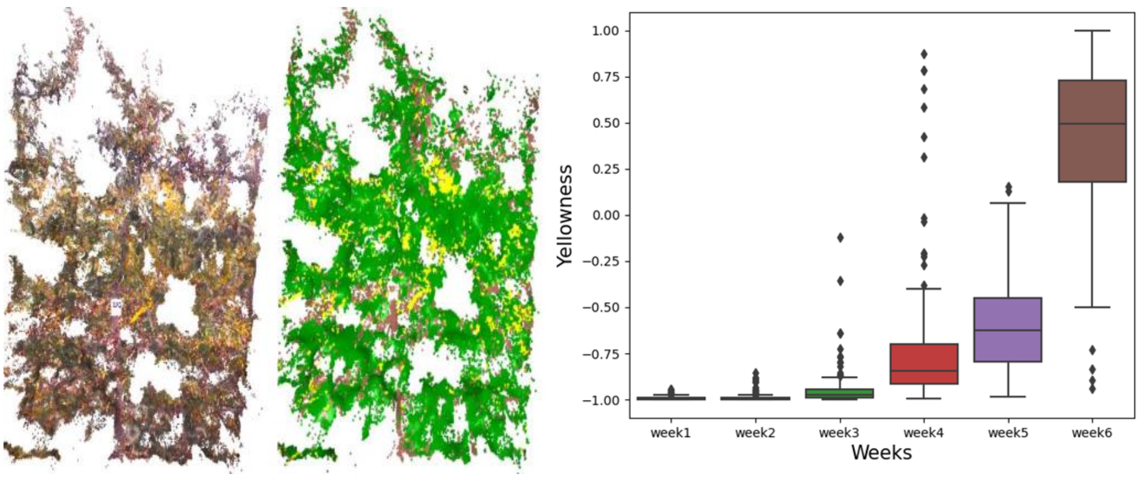

Figure 11: (Left) Segmented point cloud and clustered point cloud (i.e. Green, Yellow, and Trunk cluster) of a sample tree. (Right) Yellowness for all trees during different weeks of study.

Results & Discussion: The yellowness index of each tree was calculated over the 6 weeks of the study. Figure 11(right) shows a plot of yellowness for all trees during the 6 weeks. The results show a general trend of yellowness increasing with each week (i.e. trees turning more yellow), as expected. The boxplot shows that all trees start out at a yellowness index of ~-1 during the first week (i.e. all trees were completely green). However, at week 3 there is an increase in the yellowness index. At week 4, the change is more prominent where there is a significant increase in yellowness index. By week 6, most of the trees have a high $yellowness$ value (i.e. they are almost through the complete color change and have turned yellow). However, there are still some trees with a negative yellowness index (i.e. still on the greener side).

The trees were classified into 5 classes of N status: Very low N (N < 1.7), Low N (1.7 < N < 2), Good N (2 < N < 2.4), High N (2.4 < N < 2.6), and Very high N (N > 2.6). Figure 12 shows the yellowness values by week with a color code assigned to trees from the different N classes. Trees with lower N start the transition earlier in the season. At week 4, this is more clear as the trees with lower N start the transition to yellow. However, most of the higher N trees are still towards the greener side. At week 6, most of the lower N trees are already at yellowness index of +1 (i.e., completely yellow), however, there are still quite a few higher N trees transitioning color. This transition is affected by several factors including environmental stress, nutritional stresses, and aging. The results show correlation between the yellowness index and the N content at different weeks (R^2 = 0.14 – 0.18) and indicate that the yellowness index of a tree at different weeks and the pattern in which they are changing can be a potential indicator of the N status in the tree. Also, weeks 3-4 can be a good time to differentiate between high N and low N trees as the trees with low N start to change color.

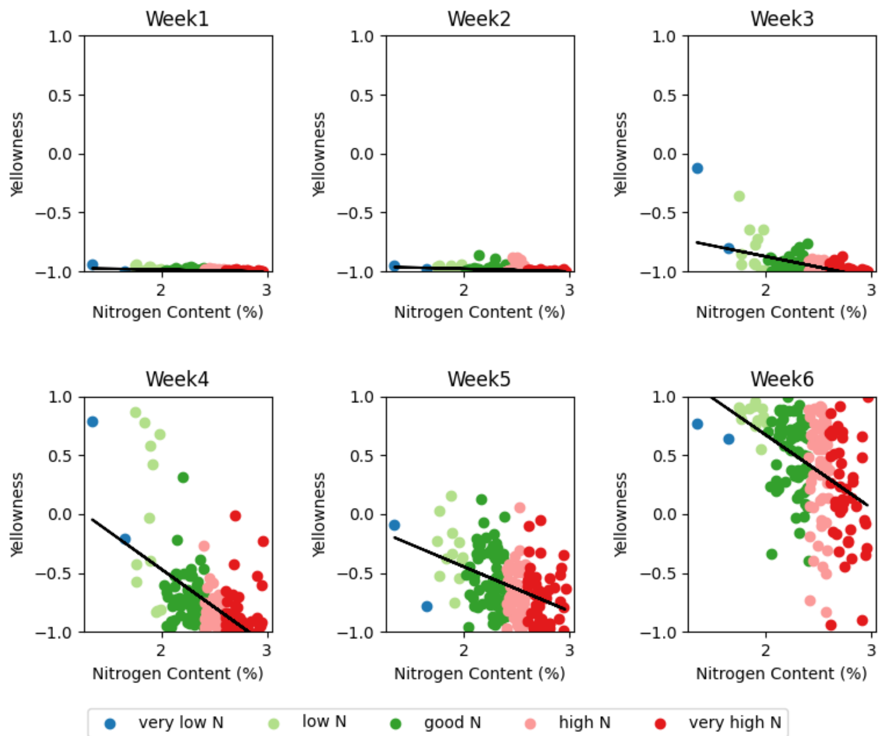

Figure 12: Yellowness during different weeks of the study for trees with different nitrogen levels. Week 1 started on October 8, 2021.

Objective 2: Decision support tool (WSU/OSU joint lead)

Methods: The parameters discussed above were used in the decision support tool to decide the fertilization rate for individual trees. A simple gradient boost regressor was used to fit a model to predict the leaf N concentration using the trunk cross-sectional area and normalized canopy area as input. The predictions were further classified into 5 classes: 5 - Very low N ($N < 1.7$), 4 - Low N ($1.7 < N < 2$), 3 - Good N ($2 < N < 2.4$), 2 - High N ($2.4 < N < 2.6$), and 1 - Very high N ($N > 2.6$) based on their predicted leaf N concentrations. The model was trained on 50% of the data from 2022 and tested on the rest of the data from 2022. A proportional quantity of fertilization was calculated based on the classification of the trees (i.e. tree with class 1 targeted with a rate of 10 lbs/acre, and a tree with class 5 targeted with a rate of 50 lbs/acre by varying the volume of liquid sprayed around the tree).

Results & Discussion: The model was able to predict the leaf N concentration with an RMSE of 0.32%. With this level of RMSE, the predicted classes of the target trees are expected to be within one class (one class higher or lower). The current model for decision support is expected to be further improved with the integration of fall leaf color, the yield from previous years, and experts' rating as feedback to the system.

Objective 3: Variable rate N application

Task 1 – Vehicle localization (OSU lead, WSU participant)

Methods: To accurately apply N autonomously based on visual tree metrics, it's essential to identify the specific tree being examined, ensuring the data corresponds to the correct tree. This was accomplished by localizing the ground robot (Clearpath Warthog) on a pre-made map of the orchard using a particle filter system. Particle filters localize by first scattering a large number of potential position estimates (i.e. the particles) across the map. Then, as more data from the environment is captured while the robot traverses, including RGB-D images of trunks and wheel odometry, the algorithm uses probability theory to update predicted positions until only one position estimate remains. For brevity, we do not present the details of the algorithm in this report. Figure 13 shows a graphical depiction of the multiple stages of the algorithm.

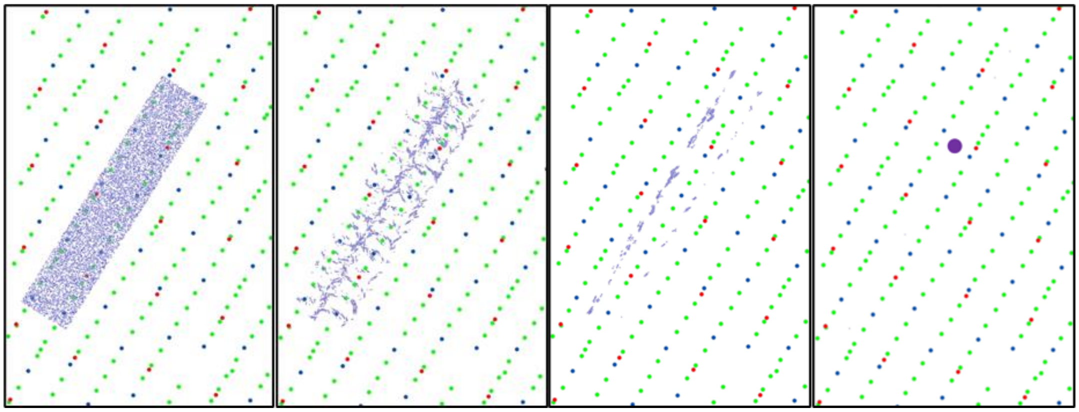

Figure 13. Graphical depiction of four stages of the particle filter system. In the map, posts are red dots, treatment trees are large blue dots, other trees are green dots, particles are small blue dots, and the final position estimate is a large purple dot. The left image shows the initial spread of the particles. The next image shows the system once a single tree has been seen. The third is after several trees have been seen. In the right image the particles have converged (purple dot).

Results & Discussion: The localization system was tested in real time in the field and shown to perform well. Localizing in an orchard is difficult primarily since everything looks very similar: the trees all look the same, are similarly spaced, and are in straight rows. The addition of the width as a metric for calculating the probability of the particles aids in reducing this ambiguity, but it's still a challenge. On the other hand, once the robot has converged on the correct general location, the error in the localization is largely a product of the accuracy of the camera system, the accuracy of the transform between the camera and the robot, and the accuracy of the map. To this end, most of our analysis focused on whether the algorithm converged correctly and how long it took to converge, for various sized initial spreads of particles.

To evaluate the system's performance, we tested 3 different sizes of spread in the field, conducting approximately 20 trials for each (see Table 1). Each trial started at a unique location on the map. For clarity, the spread height and width indicate the size of the initial spread; for instance, the example in Figure 13 has a width of 4 m and a height of 20 m. The 'distance to converge' metric is how far the robot traveled before the system converged. It's evident from our trials that there was a large variation in convergence time and distance, as indicated by the large standard deviation, but as one would expect, larger initial spreads result in longer convergence times and reduced convergence accuracy. The particle system converged to the correct location 90% of the time even when the initial spread was large.

Table 1. Results for the particle filter trials. The spread width and height give the initial spread area, the correct convergence gives the % of trials that converged correctly. The time to converge is how long it took the system to converge, and distance to converge is how far the robot traveled before the system converged.

Spread width (m)	3	4	12
Spread height (m)	10	20	20
Num. trials	21	21	32
Correct convergence (%)	95.2	90.5	90.6
Avg. time to converge (s)	16.19 +/- 10.47	18.73 +/- 9.69	25.39 +/- 8.66
Avg. distance to converge (m)	3.99 +/- 2.62	5.16 +/- 2.73	7.17 +/- 2.67

Task 2 – Precision spraying (WSU/OSU joint lead)

Methods: The N application rate recommended by the decision support system was administered through an automated system capable of self-localization within the orchard. The main platform for the integrated system was a Clearpath Warthog ground robot as shown in Fig. 14. A 25-gallon reservoir was mounted on top of the ground robot to store the liquid solution to be sprayed (the system was tested with water in field trials). The solution was then applied with an 80^o flat-fan spray nozzle (TeeJet Technologies, Illinois, USA) and 12V solenoid valve (TeeJet Technologies, Illinois, USA) that was attached to the ground robot at an angle of 20^o with the robot body, oriented towards the root zone of the trees. The solenoid valve was actuated with a microcontroller. The front-facing 2D LiDAR (Sick LMS 111-10100 (Sick AG, Waldkirch, Germany)) was used to capture a laser scan in front of the robot to control its motion and ensure safe operation. The side facing Realsense D435i (Intel Corp., CA, USA) was used to detect trees while the platform was moving. The LiDAR, Realsense D435i camera and Arduino were connected to a Jetson Nano (NVIDIA Corp., CA, USA) on the Warthog. Both the Jetson Nano and the Warthog computer were connected to a common network via an ethernet switch (TP-Link Tech. Co., China). The entire system can be accessed through the network switch or wireless router connected to the switch. The spray line was pressurized at 40 psi with a pressure regulator and relief lines. Two 12V batteries were connected to power the pump for spraying and actuating the solenoid valve, while the Warthog was powered with rechargeable lithium batteries. The integrated system was self-contained with all power required for running the equipment on-board.

Figure 14: The integrated system: a) from front and b) from back. The front-facing LiDAR and the side-facing camera are visible in a), while the spraying nozzle has been zoomed in on b). The yellow tank on top contains the spray solution.

Results and Discussion: The integrated system accurately sprayed a quantity of liquid, as decided by the decision support system, directly onto the root zone of the trees. The integrated system accomplished the following tasks: i) Autonomous navigation within a row in the orchard, ii) Detection of trees and localization in the orchard, and iii) Precise application of the designated N amount to the tree's root zone. The overall system integration was implemented using Robot Operating System (ROS). These tasks were run on multiple computers connected to the same network through the network switch and communicating with each other in real-time through ROS.

i) Autonomous Navigation Node: The autonomous navigation system was able to identify the left and right tree lines based on the incoming laser scans by fitting a least square fit line to the points on the left and right. The center line was the average of the right and left lines (Fig. 15). The heading and position of the robot were identified using the orientation of the center line. Angular velocities were controlled to maintain a constant distance of 0.1m to the left of center and oriented parallel to the orchard rows with a constant linear velocity. The system was set to observe 5m in front of the robot.

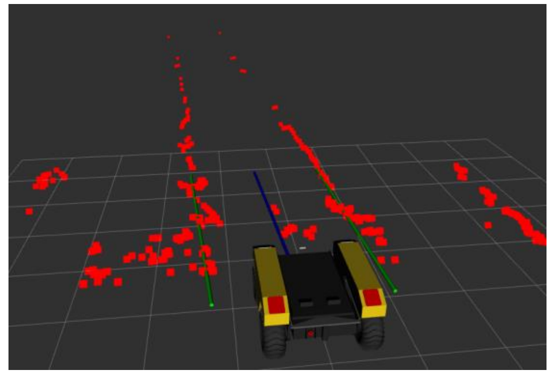

Figure 15. An instance of the laser scan (red dots) and the fitted line for left and right rows. The green lines show the fitted line for left and right tree rows, and the blue line shows the center line.

ii) Localization Node: The localization system, as discussed in Task 1 of Objective 3, was used in the integrated system to identify the position of the ground robot with respect to the trees on the map. The localization node identified the tree trunk and estimated its diameter in real time, which was used as an input for the decision support system to decide the amount of liquid to be sprayed on each treatment tree.

iii) Sprayer Node: A proportional amount of fertilizer based on the classification result from the decision support system was applied to the test trees. The decision support system took as input the real-time diameter estimation from the localization node and normalized area for the corresponding tree from the 2023 data that was calculated in early August (the system was tested in late October). The spray was actuated for a time proportional to the required volume of spray as would be required for a solution with water-soluble 20% N fertilizer diluted at 1 lb of fertilizer per liter of water. The flow from the spray nozzle was calibrated using a measuring cylinder with different volumes. The calibration model allowed for obtaining a precise volume output from the spray nozzle within 5 ml. The nozzle remained off until it reached the test tree and actuated at 5 different distances from the test tree based on the required volume, which was proportional to its N requirement. The spray time was calculated based on the flow from nozzle, the desired concentration of fertilizer in the solution, and the linear velocity of the robot to uniformly apply the solution to the test tree. The spray area of the system was tested using two 4ft wide paper rolls spread 2 ft to the left and 2 ft to the right from the tree trunk (Fig. 16). To make the spray area visible, water-soluble food coloring was added to the water.

(a) (b)

Figure 16: Paper roll spread out in front of the tree a) before spraying and b) after spraying. The red color on the paper shows the sprayed area.

Results & Discussion: The system was able to: correctly center itself in the row, autonomously navigate within the rows, precisely localize within the orchard, and spray the solution at desired locations around the test trees. The sprayer was able to target the spray within an area of 4ft of the target trees on the sides and 1.5 ft on the front of the trees, which was within the root zones for the target trees. The video of the overall system operating in the field can be found by clicking the link.

Executive Summary

Project title: Decision Support Tool for Precision Orchard Management

Key words: Precision fertilization, variable rate technology, computer vision, deep learning

Abstract: The standard practice of broad-acre orchard management does not result in targeted actions that are optimal for individual trees – this reduces the impact of management decisions and wastes resources while falling short on achieving the yield and quality potential of individual blocks. The primary goal of this project was to improve fruit quality and yields by managing the Nitrogen (N) of individual trees through a combination of automated sensing, machine learning, decision support tools, and variable rate application technology. To accomplish our goals, we created three specific research objectives:

1. Develop a ground vehicle-mounted sensor system that i) maps the geographic location of individual trees within an orchard block; and ii) measures plant parameters (e.g. shoot vigor, trunk cross-sectional area, and fall leaf color) to estimate the N status of individual trees
2. Develop a decision support tool that recommends N application levels per tree and tracks the tree's long-term response
3. Develop and demonstrate a proof-of-concept precision spray system that localizes the vehicle with the orchard map, identifies the neighboring trees, and then selectively applies the desired level of N within the root zone

Over the duration of the project we collected leaf samples from 200 treatment trees within the test block (Jazz cultivar; Yakima Valley Orchards, Posser, WA). Mineral analysis showed that most treatment trees had excess leaf N as opposed to too little and that N status remained relatively consistent year to year. To automatically detect trees and estimate trunk cross sectional area, we used a consumer grade RGB-Depth sensor and deep learning. Compared to manual measurements with calipers, the mean absolute error of the trunk width estimation algorithm was 2.56mm. The deep learning algorithm was then used in conjunction with RTK-GPS to create an accurate map of all trees and posts within the test block.

To estimate tree nutrition with non-contact visual sensing, we created new techniques that measured canopy area and temporal changes in leaf color. We found that normalized canopy area is highly correlated (negatively) with the target N application rate for the upcoming season (based on human expert guidance). The use of normalized canopy area as a measure of canopy vigor and an indicator of the N need of individual trees is promising. Also, the yellowness index of a tree at different weeks and the pattern in which they are changing can be a potential indicator of the N status in the tree. We observed that Weeks 3-4 can be a good time to differentiate between high N and low N trees as the trees with low N start to change color.

Our final contribution of the project was an in-field demonstration of autonomous variable rate application with a ground robot. Using a particle filter system that incorporated the orchard map and trunk segmentation algorithm, the robot was able to quickly and accurately localize with respect to individual trees without the need for GPS. A decision support tool incorporating a simple gradient boost regressor was used to determine the fertilization rate for individual trees. The N application rate recommended by the decision support system was applied through an automated liquid spray system. The integrated system accomplished the following tasks: i) Autonomous navigation within a row in the orchard, ii) Detection of trees and localization in the orchard, and iii) Precise application of the designated N amount to the tree's root zone. The sprayer was able to target the spray within an area of 4ft of the target trees on the sides and 1.5 ft on the front of the trees, which was within the root zones for the target trees. Future work will focus on using historical data (e.g. yields) to improve the decision support tool.