FINAL PROJECT REPORT

Project Title: Robotic scout for tree fruit

PI: Tony Koselka
Organization: Vision Robotics Corp
Telephone: (858) 523-0857, ext 1#
Email: tkoselka@visionrobotics.com
Address: 11722 Sorrento Valley Rd.
Address 2: Suite H
City: San Diego
State/Zip: CA 92121

Cooperators: None

Other funding sources

Agency Name: USDA, SCRI Funding
Amount requested/awarded: $200,000
Notes: This is matching funding per the CASC SCRI project

Total Project Funding:

Budget History:

Item	2010
Salaries	$111,707
Benefits	$44,056
Wages	
Benefits	
Equipment	
Supplies	
Travel	$11,000
Miscellaneous	
Scout Prototype	$8,237
Subcontract to CMU (field expenses for integration)	$25,000
Total	$200,000

Footnotes: In addition to the above budget, in 2010 project included $2875 in WTFC Collaborative expenses.

OBJECTIVES

Vision Robotics Corporation (VRC) has been working closely with the Washington Tree Fruit Research Commission (WTFRC, or the Commission) to develop a robotic harvester for tree fruit. The first step is the development of a Scout that scans the trees with cameras, and locates, counts and sizes the fruit. The second step is the development of a Harvester that uses the location information from the Scout to cost-effectively pick the fruit. This project is part of the "Comprehensive Automation for Specialty Crops" (CASC) SCRI project.

The goals for 2010 were to improve the apple detection and sizing accuracy for both red and green fruit, and to upgrade the prototype Scout's robustness and operating scale. The project was to culminate with field tests to demonstrate crop load estimation on a significantly larger scale than in previous years. The specific objectives included:

1. Crop Load Estimates
    a. Obtain 90%+ yield accuracy for less densely packed orchards;
    b. Obtain 85%+ yield accuracy for highly clustered trees;
    c. Benchmark 2009 sizing performance.
2. Prototype Upgrades
    a. Operate in temperatures in excess of 100°F;
    b. Incorporate industrial flashes;
    c. Upgrade mast to be lighter, more robust and easier to handle.
3. Increase operational speed to between 1 and 2 mph.
4. GPS-reference scans and feed data into the Carnegie Mellon University (CMU) GIS database.
5. Increase analysis speed to where data from a 100' scan of less densely packed fruit can be completed within 30 minutes.
6. Integrate with the CMU Autonomous Prime Mover (APM).

VRC originally proposed a test plan to scout different varieties, growing configurations and stages of growth. Specifically, the field tests were to consist of three separate trips to Washington orchards to scout six different 1 acre sub-blocks, one of which was to be scanned during all three trips. The three trips were to permit testing at three different growing stages: one week before harvest, four weeks before harvest and when the apples are approximately one inch in diameter. Based on input from the Commission, the test plan was modified to scouting two blocks where each was at least five acres in size. These tests would result in significantly more data collected than originally planned. As such, both sets of field tests were conducted during the same trip to Washington.

In addition, the Scout system was designed to work with fruiting walls approximately 18"-24" thick. One of the blocks scanned during the field tests had trees in excess of 36"-48" thick. (This block was selected for the strength of the red coloring of its apples at the time of scanning.) Apples on the far side of the tree were not visible to the Scout cameras, so a statistical analysis is required in order to estimate the crop load.

VRC originally requested $275,000 dollars for the project and $200,000 was awarded. The decrease in budget was partially offset by decreasing the number of field tests and trips to Washington. In addition, less effort was available to refine detection and sizing algorithms, and to work with smaller apples.

SIGNIFICANT FINDINGS

Field Test Results

McDougal Farms, Ambrosia Block (Green Fruit)

- The crop load estimate for 6 acres was 560,358 which was
- 75% of the load extrapolated from hand count of 240' to estimate of 742,388.
- Raw count was 299,997.
- Counts from four trials each of two 60' hand count sections each self-consistent within 5%.
- Median apple diameter estimate was 2.60” which was within 1% of the median of 2.58” of the 240' of hand measured apples.
- Raw median apple diameter was 2.62”.

Washington Fruit and Produce, Gala Block (Red Fruit)

- The crop load estimate for 11 acres (every other row was scanned and the results were doubled) was 1,445,143 which was
- 94% of the load extrapolated from hand count of 240' to estimate of 1,532,232.
- Raw count was 284,477 for 1/2 of the block (5.5 acres).
- Counts from four trials each of four 60' hand count sections each self-consistent within 8.5%.
- Median apple diameter estimate was 2.56” which was within 3% of the median of 2.63” of the 240' of hand measured apples.
- Raw median apple diameter was 2.49”.

Crop Distribution

- The Scout scanned 11.5 acres over two blocks, which is a significant amount data collected about the orchards. While two blocks do not represent a large statistical data set, it does demonstrate the potential for the Scout to enable precision farming.
  - The 6 acres of ambrosia apples consisted of 30 rows and approximately 5 miles of trees. When broken into approximately 16' sections, the count/section values had
    - median 340,
    - standard deviation 158.7.
  - The 5.5 acres of gala apples consisted of 26 rows and approximately 4.5 miles of trees. When broken into approximately 16' sections, the count/section values had
    - median 513,
    - standard deviation 128.4.

2009 Sizing Benchmark

- As noted above, the 2010 sizing performance was extremely accurate, within 3% for the median apple size in the ground truth sections for both blocks scanned. This strong performance was predicated on benchmarking the 2009 results; analyzing actual sizing performance; improving sizing algorithms; and realizing benefits resulting from enhancements elsewhere in the Scout system.
- Vision Robotics benchmarked the 2009 software's sizing performance by scanning apples of known sizes in several different configurations.
    - The raw average apple diameter was 2.2" while the true average diameter was 2.9".
    - The data have been used to create a statistical model that can be applied to correct for such system biases.

Speed and Robustness

- Using standard off the shelf computers, industrial flashes, forced air cooling, etc. the Scout worked continuously in temperatures well in excess of 100°F during the field tests.
- Ground speeds in excess of 1 mph were achieved using a camera frame rates of 20 images per second and the new flashes.
- The industrial flashes provide sufficient light to enable the Scout to operate in all sun conditions, a flash rate high enough for production, and robustness that approaches that required for production.
- The system has been optimized such that the median time to analyze 100' of data was 30 - 70 minutes, depending on the block from which the data were collected.

RESULTS & DISCUSSION

Moving the Scout system towards production requires strong performance in counting and sizing the apples, high processing speed, high system robustness and an effective means of displaying the data. In 2010, all these components were significantly improved.

Estimation Performance

Bias Correction

In a crop load estimation system such as the Scout, it is likely that biases will inherently be present. Causes of such biases could include software factors such as a systematic tendency to undersize fruit during analysis, as well as physical factors such as high tree thickness causing some fruit to not be visible in captured images (and thus not counted). While the variability of such biases across different orchards and apple varieties is not yet known, they are likely to remain consistent through a block. Consequently, a small set of hand collected data can be used to develop a statistical model to correct the bias, transforming "raw" data into more accurate "statistically-adjusted" data. For example, by applying bias correction, an estimate is equally accurate regardless of whether the Scout consistently identifies 99% of the apples correctly with zero double counts or false positives, or consistently estimates 80% of the apple count regardless of the number of correct versus incorrect detections. Undoubtedly, the better the system is at correctly distinguishing apples, the more likely it is to have a consistent count.

In 2010, hand counts from 60' sections were used to perform bias correction. The ratios of the median hand measured diameter to the median Scout estimated diameter were 0.98 and 1.06 for the green and red apples, respectively. These near-unity diameter scaling factors demonstrate that only a small bias was present in the median size estimates, and that potentially either an extremely small or no bias correction will be required for sizing

To correct for bias in count estimates, a scaling factor was determined for each of the green and red apple scans by averaging the ratio of the hand count to the Scout count over the four 60' sections. The resulting count scaling factors were 1.87 and 2.54 for the green and red apples, respectively. The relatively large separation of these values from unity shows that, despite being rather self-consistent, a fairly large bias was present in count estimates. A brief review of images in which detected fruit were outlined with circles suggested that a large percentage of more distant fruit were not visible to the Scout's cameras because they either were obscured by closer foliage and fruit, or were insufficiently illuminated by the flashes. In general, as illustrated in the images below, the majority of the fruit detected by the Scout were on the close side of the trunk, which roughly corroborates a scaling factor of 2. One potential approach to combating such a source of bias would be to detect tree trunks, filter detected fruit to include only those on the near side of the trees, and then double the estimated count. An alternative is for the Scout to determine the thickness of the canopy and determine the depth into which is sees a large and uniform number of apples. It could then filter out all apples beyond that depth and use a scaling factor between that depth and the total canopy thickness. Both these logical refinements would lead to a more stable bias factor across different orchards.

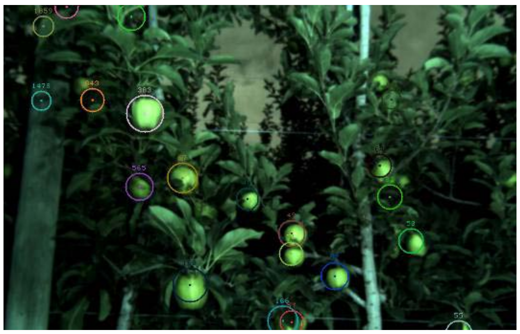

Apple Count

To improve raw estimation performance, the first step in 2010 was to further analyze the Scout's 2009 performance. The estimates of the number of apples in the 100' sections of jazz rows in the Allan Brothers orchard scanned in 2009 were typically accurate (based on 4-6 runs) to within 25% using 2009 software, with some of the runs achieving 98% accuracy. The Scout count includes correctly identified apples, doubly counted apples, missed apples and false detections. The results of these analyses, coupled with experience gained improving algorithm performance in 2009, provided target areas for enhancement in 2010.

VRC has implemented an improved detection algorithm with better identification of individual apples (particularly those of mixed color) and stronger performance in identifying individual apples within clusters. The new algorithm uses the same software for detecting both red and green apples, but requires different input filters. The visual odometry software module is used to determine how much the cameras have moved between pictures and to determine the relative positions between cameras. The module has been enhanced to better correlate the portions of the images that overlap between cameras, which helps to eliminate double counts due to multiple cameras seeing the same apple (the largest number of errors in 2009) and enables improved location and size determination by incorporating more views of the same apple from different perspectives.

Significant improvements to the collection software and the prototype design also improved detection performance. Using data collected from the 2009 field tests, the team optimized the camera/mast configuration, adding one additional camera pair (nine as opposed to eight in 2009) and changing the relative locations and orientations to better image entire trees. Additionally, camera settings were adjusted and the auto-exposure algorithms were updated. Taken together, these modifications have improved the capability of the Scout to see the apples on the trees, which leads directly to improved estimation accuracy.

The figures below show the raw and statistically-adjusted counts for the 30 rows of green apples, and the 26 rows of red apples. The median and standard deviation of the statistically-adjusted row counts were 17820 and 2799 for the green apples and 28709 and 4548 for the red apples. The plots illustrate that while a significant bias correction factor has been applied, the raw counts are overall fairly consistent from row to row. Note that row 31 of the red apple block was of substantially shorter length than the other red apple rows, giving rise to a correspondingly low count.

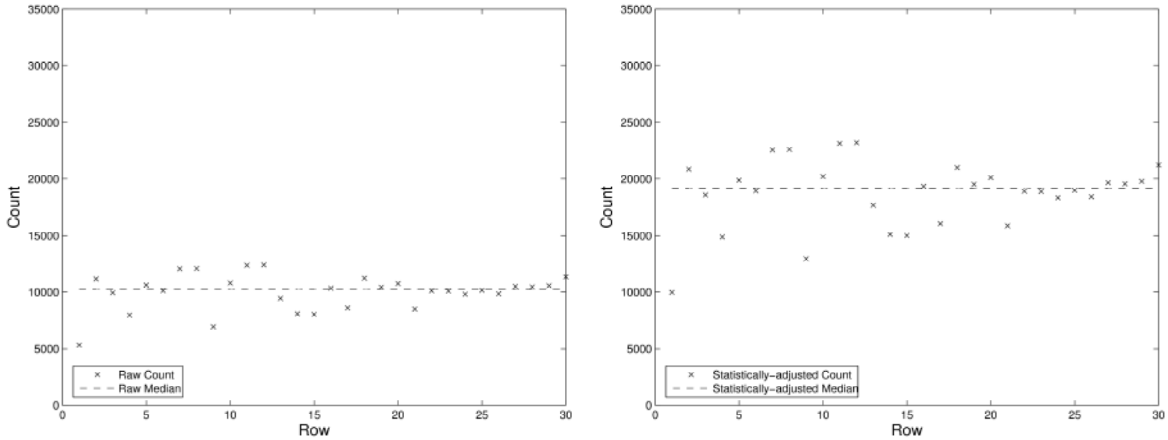

Raw counts (left) and statistically-adjusted counts (right) for 30 scanned green rows

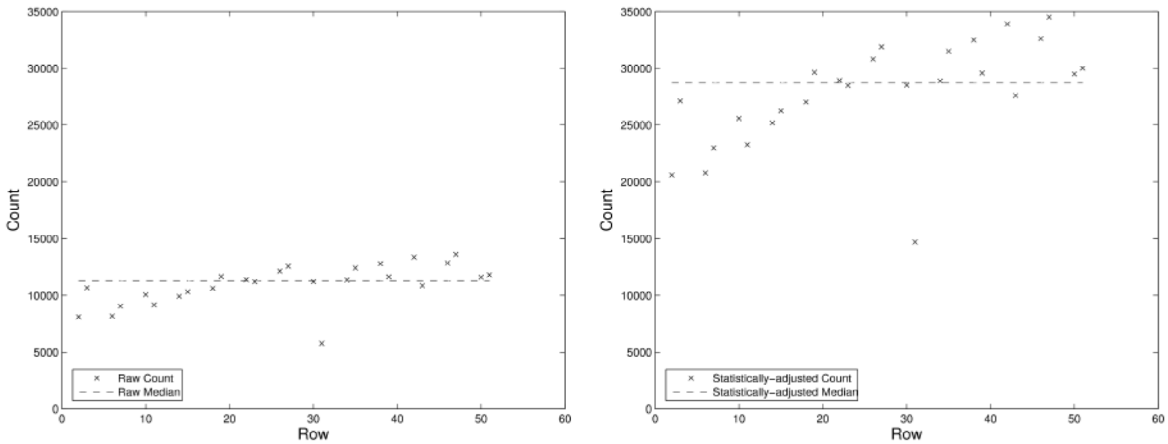

Raw counts (left) and statistically-adjusted counts (right) for 26 scanned red rows

To view the data on a finer scale, each row was divided into approximately 16' sections along the row. Histograms giving the statistically-adjusted counts per section are shown below for the green (left) and red (right) fruit. These plots demonstrate the Scout's ability to detect load variability within the block.

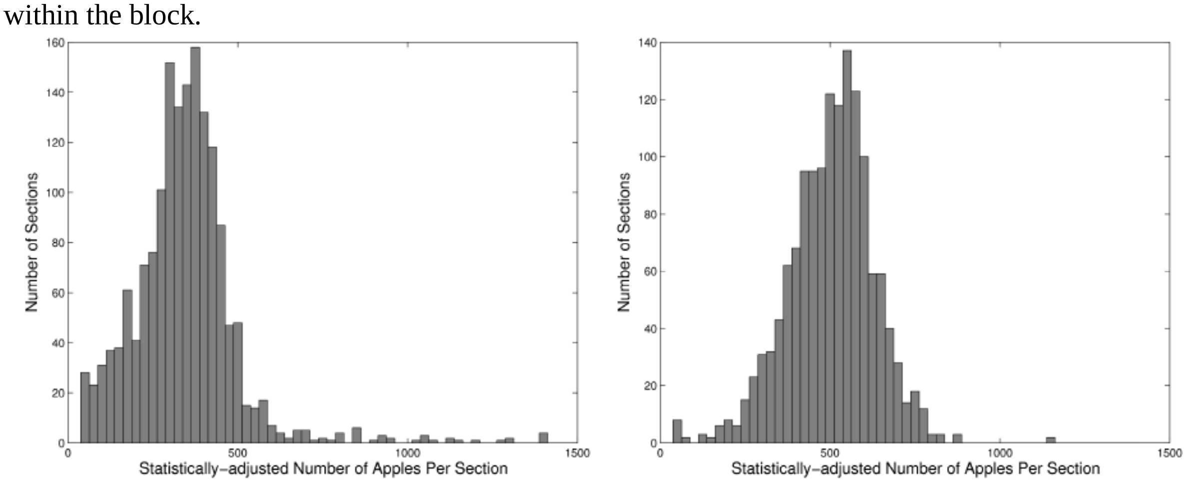

In order to develop a bias correction model and to study the consistency of the Scout's estimates, the hand counted 60' sections were scanned repeatedly and the results are shown in the table below. In all cases, the latter three trials were performed in succession, while the first trial was performed at a different time of day. The counts clearly demonstrate the consistency of the Scout estimates over these trials.

Row	Scout Counts	Consistency
Green 18	829, 780, 772 and 818	7.4%
Green 24	807, 812, 780, and 768	5.7%
Red 10	769, 817, 833 and 819	8.3%
Red19	781, 737, 744 and 726	7.6%
Red 34	777, 851, 844 and 878	12.9%
Red 43	764, 733, 724 and 704	8.5%

Sizing

Analysis of the 2009 sizing performance revealed that partially occluded apples tended to lead to size underestimation with a larger variance than true distributions; the estimates of the range of apple sizes tended to be wider, flatter and shifted to smaller sizes than the actual crop. Sources such as this introduce system biases which can be reduced through the use of statistical modeling. A statistical model was created to adjust the size distribution to addresses these expected inaccuracies. More and improved data from this year's field tests will enable refinement of the model in 2011.

As reported last year, the 2009 raw average size estimates were approximately 20% too small. For the run shown in the below, the raw average apple size was 2.2" diameter, which is 24.1% less than the hand measured average of 2.9" diameter. The histogram of the data after the statistical model was applied shows a mean size of 3.0" diameter, or 3% larger than ground truth. Thus, the statistical modeling is effective in adjusting the mean of the size distribution. The variance of the distribution, however, remains larger than the true variance. This result is, in part, due to the need for a larger sample size when developing the statistical model. Such a larger set will be available when statistical models are developed based upon field data as opposed to laboratory data.

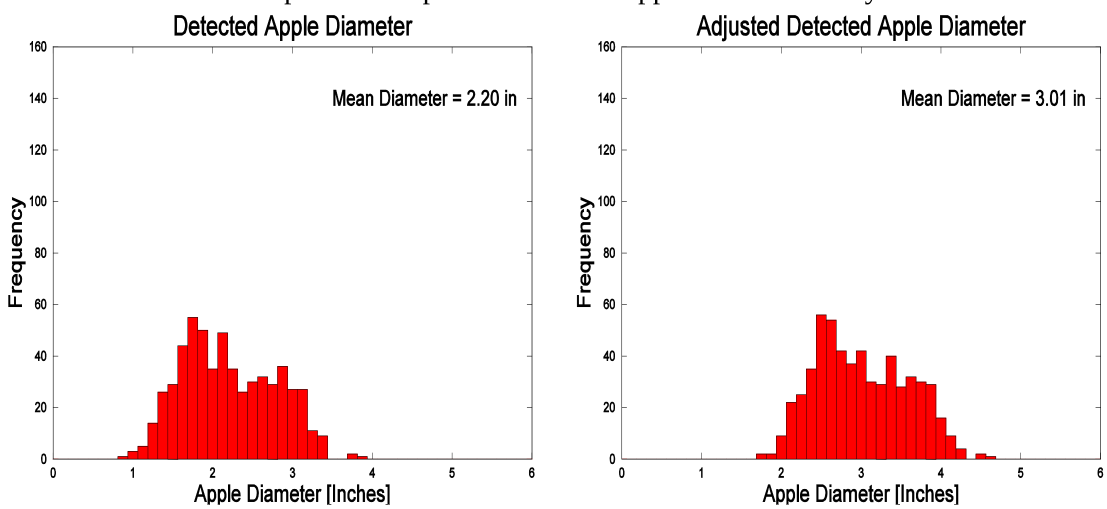

Raw (left) and statistically-adjusted (right) size distribution for a 2009 scan

When compared to 2009, the 2010 Scout produces a better distinction between an individual and the surrounding fruit, leaves and branches in the images. This delineation directly leads to the software more accurately detecting the perimeter of the apples in the images, thereby significantly improving the raw average sizing performance. The 2010 raw and statistically-adjusted (using the simple scaling model) size distributions for the 6 acres of green fruit and the 5.5 acres of red apples are shown below, as are the aggregate hand measured size distributions of fruit in the four 60' sections for green and red apples. Visually comparing the histograms for the raw and hand measured fruit diameters immediately illustrates that very little bias is present in the median size; however, the Scout estimates display a larger variance, as is expected. After bias correction with a near-unity scaling factor, the median green apple diameter was 2.60”, which was 0.78% larger than the hand measured median diameter of 2.58”. Similarly, the bias-corrected median red apple diameter was 2.56”, which was 2.66% smaller than the hand measured median diameter of 2.63”.

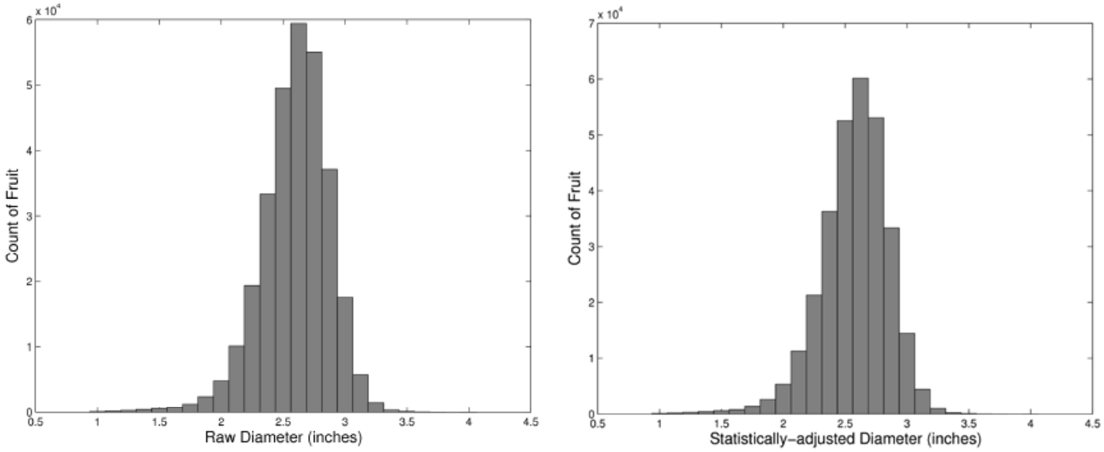

Raw size estimates (left) and statistically-adjusted size estimates (right) for 30 scanned green rows

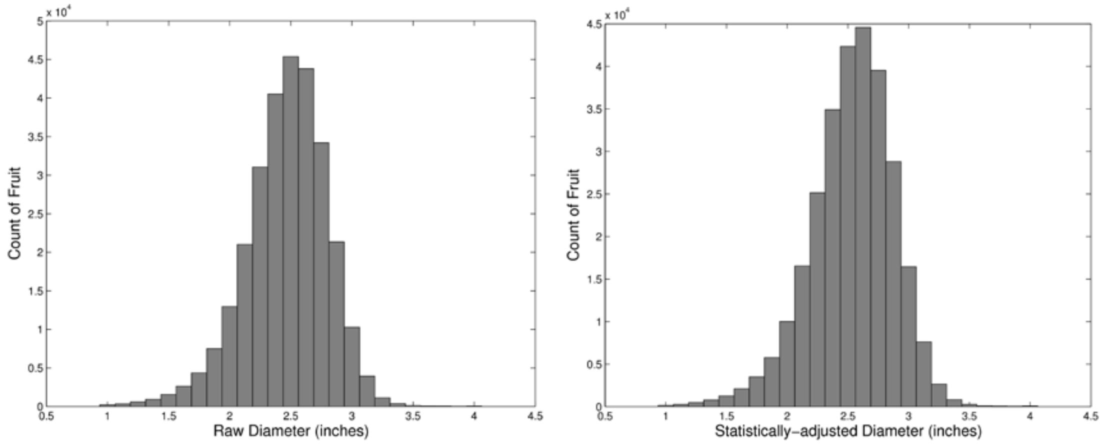

Raw size estimates (left) and statistically-adjusted size estimates (right) for 26 scanned red rows

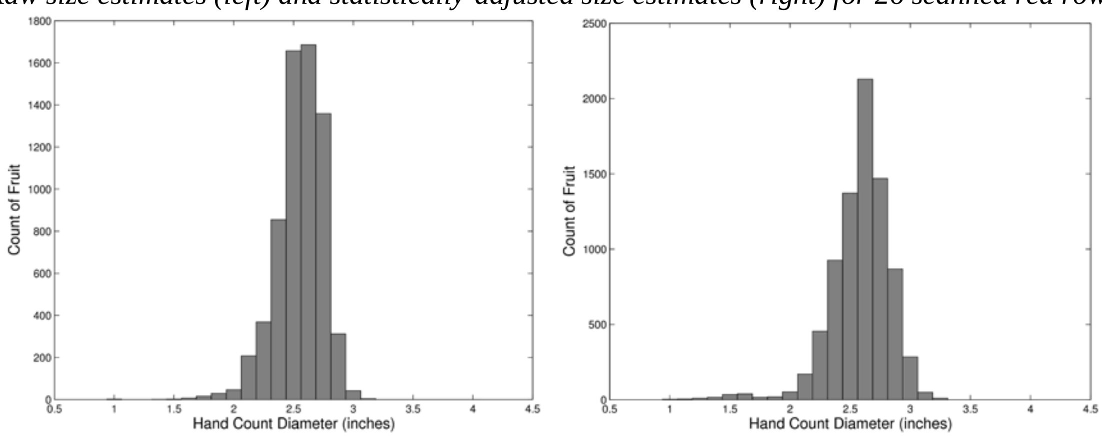

Hand measured apple sizes for 240' of green (left) and red (right) apple

System Robustness

The 2010 Scout prototype and camera module, shown below, represent significant upgrades towards a production design, but the resources necessary to fully weatherize them were not expended. For example, each camera pair is now in a closed module that is straightforward to fully seal, but the time and expense were not taken to use IP65 connectors and gaskets. Similarly, the computers are more robust than those used in 2009, but they are still standard desktop models. Fully weatherized and robust computers are available, but only represent a marginal robustness improvement that, as expected, was not required this year. Active cooling of the electronics cabinet through a fan system was included in the 2010 prototype. VRC conducted several local field tests to debug the prototype and the final unit operated virtually flawlessly during the week of field tests in Washington where the temperatures were in excess of 95°F every day and above 100°F a couple of days.


Scout prototype (left) and camera module (right)

As noted, the improved robustness was a part of the requirement to ensure that the Scout can operate at a production scale. Additional improvements introduced to achieve this goal include:

- Increased Scout scanning speed to 1 mph;
    - Increased camera frame rate;
    - Decreased image density (pictures per inch).
- Incorporated GPS system to geo-reference data.
- Decreased number and increased robustness of electrical connections.
- Debugged software to eliminate crashes.
- Incorporated industrial flashes.

<u>Analysis Speed</u>

The current Scout is approximately 20 times faster than in 2009 when analyzing data from 2009, meeting the goal of analyzing 100' of data within 30 minutes. The speed gains were achieved primarily through parallelization and decreased analysis time because of the new detection algorithms. Analysis times for 2010 data are somewhat longer, with median times of 70 and 59 minutes for 100' of green and red apples, respectively. This increase can be attributed primarily to the blocks, which have significantly more fruit than those scanned in 2009, and to the fact that the scans were significantly longer (the longer runs require the software to track more fruit during each run).

<u>Integration with the APM and CASC GIS System</u>

In 2010, the APM towed the Scout throughout the field tests at the Washington Fruit and Produce orchard. VRC and the CMU team spent almost a full day integrating the two systems, primarily updating the APM software to correctly turn between the rows when towing the Scout. The two robots completed the red apple scans over the course of the next day and a half.

VRC will provide the crop load estimate data to CMU shortly for integration into the GIS database. The data transfer is expected to occur without incident and results are expected soon.

<u>Data Visualization</u>

VRC has created a framework for viewing data output by the Scout to provide detail and a debugging environment at VRC. The crop load data can be overlayed onto a Google Earth map of the block to show the crop load and sizes for various resolutions. The yield and median size maps for the red apple data broken down into approximately 16 foot sections along each row are shown below, as are the corresponding maps for the green apple data. In each case, red indicates lower counts (smaller sizes), yellow indicates medium counts (sizes) and green indicates higher counts (larger sizes). Note that some variability is present due to inaccuracies in raw received GPS data. Such inaccuracies likely account for instances where data which should truly appear in rows which are quite red being shifted to appear in neighboring rows (making them very green). Data for a single row can be shown at any resolution, approximately 16' sections here, with darker indicating higher counts. The size histogram for any section can be displayed by clicking on a section of interest, as also shown below.

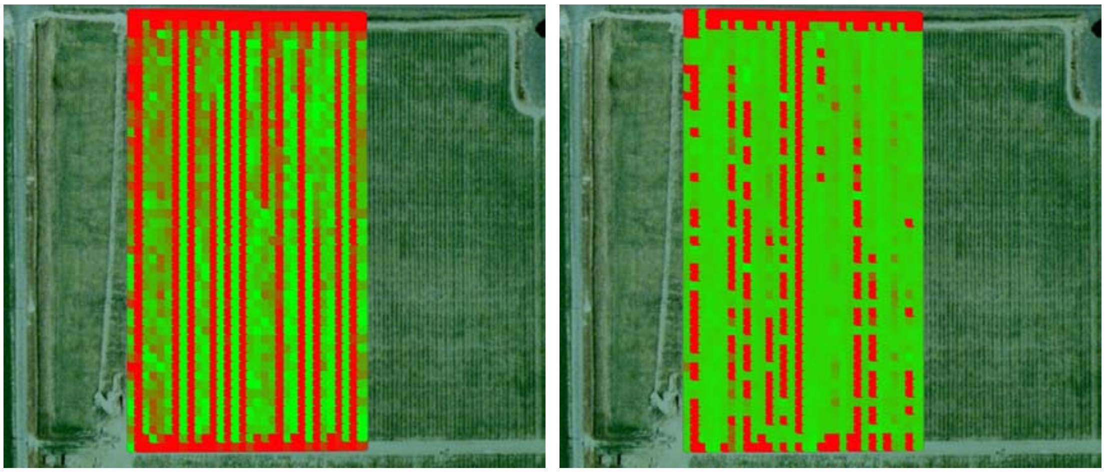

Yield map (left) and median size map (right) for the 26 scanned red rows

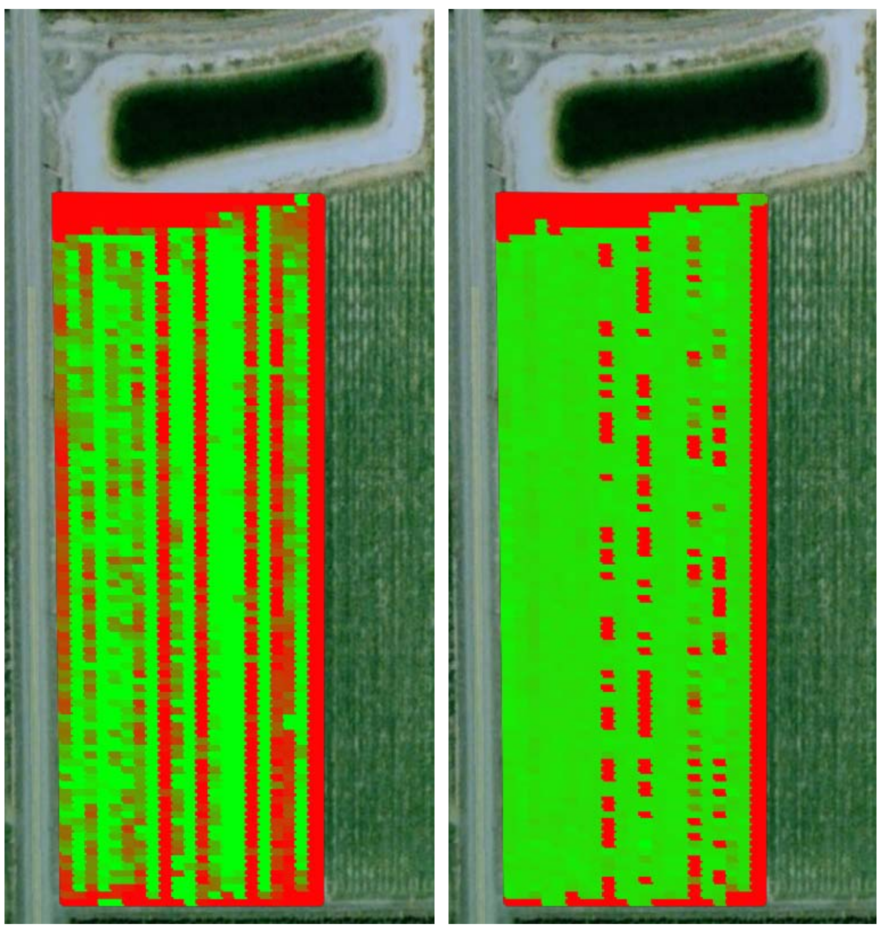

Yield map (left) and median size map (right) for the 30 scanned green rows

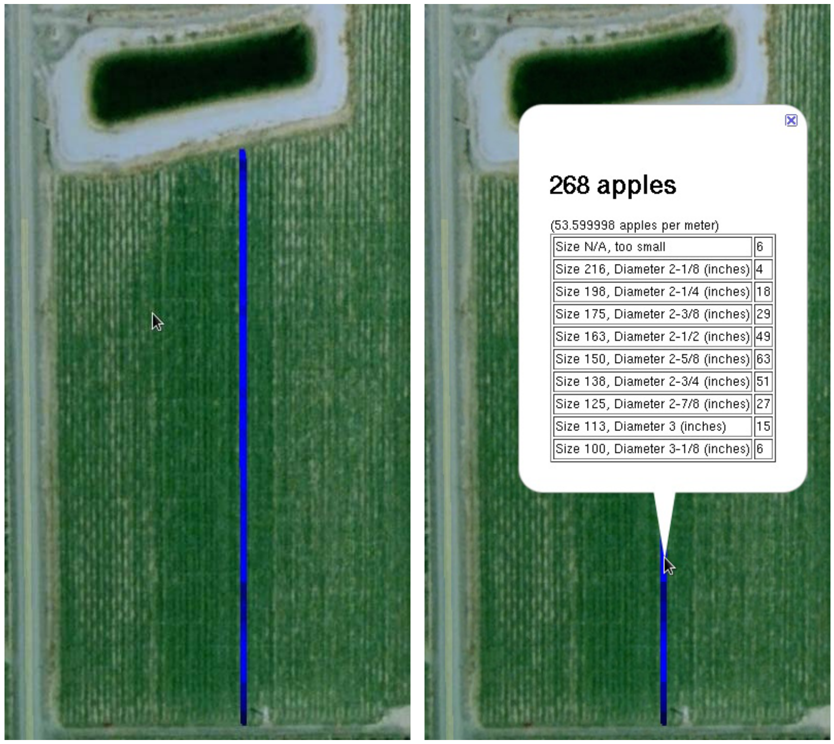

Yield map for a single row (left) and size information for a 16' section (right) for a green apple scan

Field Tests

The team wishes to acknowledge and thank everyone that helped make our field tests a success. This includes the McDougals and Washington Fruit and Produce who let us into their orchards as well as helped keep us dry during those hot days. Similarly the CMU and WTFRC teams went way above and beyond reasonable effort. Collectively, we worked from before dawn to late into the night, and even through thunderstorms. Finally, we appreciate and thank the Commission for collecting the ground truth data, both the estimates for the full blocks and the hand counts of small sections within the blocks. During the tests, the Scout collected approximately 9 terabytes of data for analysis. Despite the temperature, rain and sprinkling, it operated all three days and collected data without a failure except for a couple of hard disk related crashes with one of the 12 disks used.

The Future

VRC is pleased with the 2010 progress; the detection software performance (particularly with respect to sizing) and speed was improved, and the Scout demonstrated its ability to scan large blocks with high consistency within a block. The goals for the future include a plan for continued refinement of the apple detection and sizing performance, and further increasing the processing speed with the ultimate goal of achieving real time. One specific goal for 2011 is to analyze the data collected this year to determine a statistical scanning plan to create accurate crop load estimates while scanning only portions of the orchards. An additional key goal is to collect data from a larger and more diverse set of blocks (in terms of varieties and tree configurations) to analyze the variability in the statistical models used for bias correction.

EXECUTIVE SUMMARY

The Scout scans apple trees with cameras in order to estimate crop load, and is the first step in the development of a robotic harvester for fresh apples. Collectively with the other aspects of the "Comprehensive Automation for Specialty Crops" SCRI project, the Scout will help mechanize the growing of tree fruit. The goals for 2010 were to improve the apple detection and sizing accuracy for both red and green fruit, and to upgrade the prototype Scout's robustness and operating scale.

The 2010 test results demonstrated significantly improved fruit detection and repeatability of performance. The field tests also showed that trees with lush canopies and a heavy crop load inhibit the Scout cameras from seeing the all the apples on the trees, biasing the scout estimates to be low. However, the crop load estimates for sections with verified hand counts were consistent and repeatable to within 10% across multiple trials. These results strongly indicate that the inclusion of a bias correction method will yield an accurate crop load estimate.

During the field tests, the Scout scanned 11.5 acres in two blocks, one with green apples and one where the apples had started to turn red. In the green block, the Scout counted 299,997 apples. The count after the bias correction was 560,358. In addition, the Commission hand counted apples in four 60' sections within the block. For those sections, the ratios of the hand counts to the Scout's count were 1.95, 1.71, 1.88 and 1.94, demonstrating the Scout's consistency. The raw median diameter estimated by the scout over the full green apple block was 2.62", while the median hand measured diameter over the 60' sections was 2.58", indicating that very little bias was present in the Scout's median size estimates. The scouting procedure was the same in the red apple block, and the results were similar. These results indicate that the Scout is capable of producing an accurate crop load estimate.

While providing absolute crop load estimates is essential, for the first time the Scout was able to show the apple distribution throughout a block. For the green apple block, breaking the 6 acres into 16' sections showed that the apple counts in the sections had a median and standard deviation of 340 and 158.7 respectively. Ultimately, data such as these will help decrease the growing costs for tree fruit by enabling precision farming to better target efforts and costs only where they provide benefit.

Throughout the field tests, the Scout operated almost flawlessly, including operation in temperatures well in excess of 100°F, with no failures caused by heat or continuous operation for the computers, cameras or flash system. Also in 2010, VRC increased analysis speed by approximately 20 times, but still requires an additional processing speed increase of another 20-40 times to achieve real-time operation.

The next phase of development will move the Scout towards production. One key aspect of a viable product is a simple bias correction process. Growers cannot be expected to hand count small sections of each block to be scanned. Improving scouting performance will decrease the bias correction required. One technique that should significantly decrease the variability of the required bias correction across different orchards, tree configurations and apple varieties is for the system to determine how far it can thoroughly see apples in the trees being scanned, limit counting to that depth and directly scale the count proportional to the ratio of tree thickness to visible depth. A second key requirement for the Scout is to operate quickly and cost effectively in full size production blocks. This objective can be met by statistical sub-sampling the blocks. Thus, the plan for 2011 is to quantify and implement both bias correction and sub-sampling methods. Finally, in 2011, the Scout performance will be enhanced to accurately operate in a wide range of apple varieties and orchard configurations.