**FINAL PROJECT REPORT**

**Project Title**: Robotic Scout for Tree Fruit, Phase 2


**PI**: Tim McConnell
**Organization**: Vision Robotics Corp
**Telephone:** (858) 523-0857, ext 5#
**Email**: tmcconnell@visionrobotics.com
**Address**: 11722 Sorrento Valley Rd.
**Address 2**: Suite H
**City**: San Diego
**State/Zip**: CA 92121

~~Co-PI(2):~~
~~Organization:~~
~~Telephone:~~
~~Email:~~
~~Address:~~
~~Address 2:~~
~~City:~~
~~State/Zip:~~

~~Co-PI(3):~~
~~Organization:~~
~~Telephone:~~
~~Email:~~
~~Address:~~
~~Address 2:~~
~~City:~~
~~State/Zip:~~

~~Co-PI(4):~~
~~Organization:~~
~~Telephone:~~
~~Email:~~
~~Address:~~
~~Address 2:~~
~~City:~~
~~State/Zip:~~


Formatted: Font: Not Bold

Formatted: Font: Not Bold

Formatted: Font: Not Bold

**Cooperators**:

**Other funding Sources**

**Agency Name**:
**Amount requested/awarded**: ~~$360,500~~

**Notes**:

**Total Project Funding**:

**Budget History:**

| **Item** | **Year 1:** | **Year 2:** | **Year 3:** |
|---|---|---|---|
| **Salaries** | $251,800 | | |
| **Benefits** | $75,200 | | |
| **Wages** | | | |
| **Benefits** | | | |
| **Equipment** | | | |
| **Supplies** | $1000 | | |
| **Travel** | $7000 | | |
| **Scout Prototype** | $20,000 | | |
| | | | |
| | | | |
| **Miscellaneous** | $5500 | | |
| **Total** | **$360,500** | | |

Formatted: Font: Bold

## Objectives

The objective of this project is to create a *Robotic Scout for Tree Fruit*, which is the first phase in the development of a robotic system for mechanization of growing and harvesting fresh fruit trees. The Scout scans fruit trees to determine the total crop yield, and the size, location and color of each piece of fruit in medium and high density orchards. This data is then used by the picker robot to enable it to reach critical speeds and efficiency. The objective for 2008 was to demonstrate continuous scouting of apple trees of any fruit color on real terrain.

Vision Robotics Corporation (VRC) was to demonstrate this by building a second generation prototype that will scan many trees in a row from one side. After scanning the trees, the prototype determined the location and size of the apples in view. The Scout was demonstrated to growers in orchards during the harvest season.

Specifically, the prototype consisted of a self-propelled trailer with six stereo camera pairs mounted at equal spacing on a 10 foot camera mast. The trailer pulls itself down the row, via a mounted winch, capturing color images with the camera pairs at a constant frame rate of 6 frames per second on each camera. After completion of the scan, the system analyzes the images, assisted with wheel odometry data and tilt sensor data to detect apples on the tree.

The image capturing parameters, including lighting, exposure and filters, were to be optimized to increase the distinction between fruit and non-fruit in the images. Likewise, the analysis program was to be modified with the introduction of a statistical classifier to better identify the fruit in the images. This will enable the identification of green and yellow fruit, in addition to improving current identification of ripe red fruit.

The analysis software was to be researched and developed to provide a good method of detecting green and yellow fruit.

Formatted: Font: 11 pt

## Significant Findings

Formatted: Font: 11 pt

### Mechanical System

- The Scout hardware fabrication proceeded as expected, with no surprises.
- Because the 2008 Scout prototype is self-propelled, safety mechanisms were required to ensure easy stopping (big red buttons).
- The Scout prototype was tested early at VRC for motion and sensor functionally. Preliminary data scans were taken in California's San Joaquin Valley to test grove performance and capture data for analysis. The hardware performed smoothly in this environment.
- The Scout prototype was tested in multiple orchards in Washington, scanning fruit ranging in color from very red (Jazz) to green (Granny Smith). This included scans in angled-V orchards.
- Scanning on hilly terrain in the Wenatchee area required some human intervention to prevent over-rolling on the hills.

### Software

- Continuous data capture with six camera pairs required enhancements to the USB driver software on the host computer, for more efficient buffer management.
- Frame rates of 6 frames per second proved reasonable for capturing adequate scan data.
- The tilt sensor data proved less important than anticipated, as the 2008 Scout is significantly more stable in the lateral direction than the 2007 Scout, meaning roll values were consistently very small.
- To detect green fruit, it was determined that the best strategy involved Boosting methods, combining a set of weak classifiers to create a strong classifier.
- The boosting methods required significant development, including the hand-development of training sets of known apples and non-apples and tuning of filters.
- Lighting on the mast was able to provide very good data from night scans, but had no effect on daytime imaging.

- The analysis software was able to detect nearly all apples in the images reviewed; however, false positives were also detected. It is expected that these can be greatly reduced by tracking fruit across multiple images. This work is still proceeding and results will be presented on completion, including numerical performance numbers.

## Methods

VRC developed the second generation prototype of the Apple Scout shown in the image below. The four-wheeled Scout towed itself down a row between two rows of trees. Six camera pairsthe same used for the 2007 Scout, are mounted on the mast to scan a row of trees on one side of the Scout.

A DC powered winch is mounted on the under-carriage of the trailer. The winch cable is anchored down the row (to the 4-wheeler shown in the image) and the winch retracts the cable, pulling the Scout down the row.

The electronics are powered by a portable generator (http://www.hondapowerequipment.com/products/modeldetail.aspx?page=modeldetail&section=P2GG&modelname=EU2000i&modelid=EU2000IAN) and are housed in a sealed cabinet for improved weather resistance over the 2007 Scout.

Formatted: Font: 11 pt

The *Scout* was towed at a fairly constant rate of up to about .5 mph or .75 feet per second. The cameras captured and recorded image data at 6 frames per second equating to an image about every 2 inches.

Wheel encoders accurately recorded the motion of the *Scout* in the scanning direction, while tilt sensors mounted on the mast recorded side-to-side motion, due to roll from uneven terrain. The data from these sensors is combined with vision analysis to accurately calculate the camera position at each moment during the scan.

The cameras captured each image with a constant exposure time. Artificial lighting was mounted to the boom for nighttime scans.

The recorded scan data is currently being analyzed by an enhanced version of the "Fruit Vision" processing software. As in the 2007 project, the software searches the images for fruit, however, new software has been developed to detect green fruit, using boosting techniques, followed by geometric analysis similar to that employed in 2007.

Boosting is loosely related to the concept of neural nets, and attempts to mimic the behavior of the brain. The system, like the brain, evaluates a very large number of features such as edges, smoothness and texture measures, color, etc. Via training, we know that apples have round edges, are smooth on the inside, have a certain "3-D" look, might have a stem at the top or a hole at the bottom, etc. No single feature by itself can identify the object. But a combination of a great number of such filters, weighted by importance, can. Thus we can produce a "strong" (good) classifier out of many "weak" (barely better than 50%) classifers. Boosting refers to a specific class of algorithms that examines the training data and picks (a) the feature detectors and (b) the weights.

The hardest aspects of detecting green fruit are occlusions, shadows, and shape coincidences. Apples can be partially hidden by leaves, stems, branches, tree trunks, and even other apples. We must have an algorithm that is resistant to large parts of the apple being covered up. Shadows projected on the apple can give rise to spurious edges which distract from the rounded outline of the apple. Our classifer must be able to ignore those edges, but still make use of real edges.

Because we want the apple detector to be insensitive to occlusions and shadows, we must train it on a large set of pictures of whole or partially obscured apples, some in shadows and some well lit. We must also train it to avoid classifying as an apple things that might be mistaken as apples, for example, round leaves, round patches of sky, and random collections of leaves, sky, tree, ground, etc. The result of all this training, which can take days to weeks, is a set of simple feature detectors, a set of weights, and a threshold: if you apply the feature detectors to windows in an image, multiply them by their weights, add up the results and compare the total to the threshold, windows where the total is greater than the threshold will be said to contain an apple, and windows where the total is less than the threshold will be called "non-apples." This test is performed at every possible window position in an image and at a range of reasonable scale factors.

Examples of apple and non-apple training images are shown below:

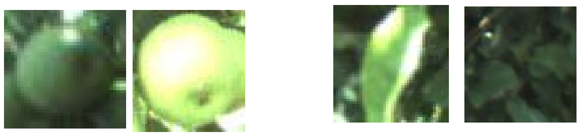

Applying a boosted classifier to an image is fast if there are not many feature detectors in the classifer, and slow if there are hundreds. In order to have good performance, a “cascade” of boosted classifiers is constructed. The cascade can also be used to filter out false positives (as dicussed further in the Results and Discussion section).

Our results indicate that our first cascade of classifiers can detect almost every green apple in an image that is not highly occluded

A good (but rather technical) introduction to Boosting can be found in **Freund, Schapire**, *A Short Introduction to Boosting*, 1999.

The dimensions of an apple can be estimated from the dimensions of the window that is used to fit the apple, as well as our odometry calculations. In a typical training image, the diameter of the apple is about 85% of the width of the image. Thus it is easy to calculate the width of the apple as it would appear right on the face of the camera. We use our geometry software to calculate how far away the apple is from the camera. Thus we can calculate how big the apple is at the tree.

Hand-measurements of exact fruit locations and sizes were taken in orchards and will be compared to the analyzed data to determine detection performance.

## Results and Discussion

**Formatted:** Font: 11 pt

VRC built and tested a prototype *Scout* to analyze trees from one side. The system consisted of a small agricultural trailer, self-propelled via a winch installed on the platform. A 10 foot tall scanning mast, on which were mounted 6 color stereo camera pairs, captured images at constant frame rate during horizontal travel. The angle of the mast can be varied to match the growing angle of the trees. The system was used to scan trees in several different Washington orchard styles, and was able to traverse the rows without flaw and capture images in many lighting conditions.

A sample of a captured image, followed by a processed image is shown below.

The processed image is displayed in gray scale, with pixels estimated to be within an apple highlighted in green. Nearly all of the apples in the image are detected by the software, along with a number of false positives. We believe these false positives will be largely eliminated by additional post-processing steps.

- Project milestones are shown in the table below.

| Milestone | Original Schedule | Final Status | Comments |
|---|---|---|---|
| Requirements specification delivered | 2/2008 | Completed | |
| Mechanical platform with mast | 4/2008 | Completed | |
| Working simulation | 5/2008 | Completed | |
| Continuous motion and data streaming | 6/2008 | Completed | |
| Orchard tests in CA | 7/2008 | Completed | |
| Orchard tests in Washington | 9/2008 | Completed | |
| Boosting methods completed | 9/2008 | Completed | |
| Geometric tracking analysis | 11/2008 | In progress | Expected completion in 12/2008 |
| Analysis completed and final results | 11/2008 | In progress | Expected completion in 12/2008 |
| Final Documentation | 11/2008 | In progrees | Expected completion in 12/2008 |

In summary, the methods employed in the development and testing of the 2008 Apple Scout have supported the hypotheses presented in the project plan. We are able to scan trees continuously, using an autonomous vehicle. The data generated in these scans contains sufficient detail and clarity for visual analysis. The analysis is able to detect the fruit in the images.

The continued geometric analysis will proceed in December to confirm that we are able to accurately determine the positions and sizes of the fruit from the scans.

This analysis will position the team well for the continued development in 2009, as laid out in the “Comprehensive Automation for Specialty Crops” project started under the USDA Specialty Crops Research Initiative.