

Results with Autonomous Vehicles Operating in Specialty Crops

Marcel Bergerman, *Senior Member, IEEE*, Sanjiv Singh, *Senior Member, IEEE*, Bradley Hamner

Abstract—Specialty crops constitute a \$45 billion/year industry. As opposed to crops such as wheat, cotton, corn and soybean, they are characterized by the need for intensive cultivation. Specialty crops growers currently face serious labor cost and availability problems, and few technological solutions exist to increase their efficiency given the past history of abundant supply of low-cost labor. This leads to an opportunity to use recent technological advances to not only increase efficiency and reduce labor costs in specialty crops production but also to support a domestic engineering solutions industry for specialty crops. We envision a family of reconfigurable vehicles that can be rapidly tasked to automate or augment pruning, thinning, harvesting, mowing, spraying, etc. They would share a common sensing and computing infrastructure, allowing applications created for one to be easily transferable to others—much like software applications today are transferable from one computer to another. In this paper we describe our work over the last three years designing and deploying a family of such vehicles, the Autonomous Prime Movers (APMs). The five vehicles completed so far have traveled autonomously over 300 km in research and commercial tree fruit orchards; preliminary results in time trials conducted by extension educators indicate efficiency gains of up to 58%.

I. INTRODUCTION

SPECIALTY crops (fruits, vegetables, horticulture and floriculture) constitute a \$45 billion/year industry. As opposed to crops such as wheat, cotton, corn and soybean, they are characterized by the need for intensive cultivation. The specialty crops industries face serious challenges today. Labor costs have increased from 38% of the net value of the farm economy to 58% in the past decade. Not only is the cost of labor limiting economic returns, but also the seasonal availability and training requirements adversely affect the cost of production. Across the board, specialty crops are facing a crisis of increasing labor costs and shortages of available labor. In addition to labor costs, an increasing consumer demand for a safe, affordable, traceable, and high quality food supply, and the need to minimize the environmental footprint, represent key challenges for specialty crop sustainability in the United States. Given the past history of abundant supply of low-cost labor, few technological solutions to increase production efficiency have been viable. The current climate gives us an opportunity to use recent technological advances to not only

This work was supported in part by the U.S. Department of Agriculture under Grant 2008-51180-04876.

All authors are with the Robotics Institute at Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. Corresponding author: Marcel Bergerman, phone: 412-513-6110, email: marcel@cmu.edu.

increase efficiency in production of specialty crops but also to support a domestic industry in engineering solutions for specialty crops.

While specialty crops as a whole represent a significant market, the needs of specific industries are diverse and there is a lack of common tools. This is as compared to automation for commodity crops for which many operations can be conducted by automating coverage patterns [11]. When the need has been substantial, growers of specialty crops have produced customized tools for specific applications. While some customization is unavoidable, few common building blocks exist that can be used to readily improve efficiency without complete systems development.

One place to look for efficiencies is in the set of tasks such as pruning, thinning, tree training, harvesting, mowing and spraying that must be conducted by vehicles moving at low speeds (typically < 3 km/h) through orchards, vineyards and groves. The state of the art today has been in the development of implements that can be attached to tractors, most of which have been designed for tasks that require much more power. Today, the need is for a generation of nimble, low-cost vehicles that can be easily reconfigured for tasks that require precision rather than brawn. Such vehicles would be up to date with standards in information technology and would fit into the electronic workflow management of crops. The vehicles would reduce the environmental footprint of production, and, most importantly would be both reliable and low-cost taking advantage of large economies of scale.

Consider, for example, the recent study by Baugher [2], who quantified the benefit of mobile platforms in orchards:

Time trials with the moveable platform and ladders were set up so they could be statistically compared. Work efficiency with the platform increased by an average of 13% for tree training, 36% for peach thinning, 50% for apple thinning, 34% for peach pruning, 53% for apple pruning, and 59% for peach harvest—all very rigorous tasks. Work quality with the platform was similar to that with ladders, and workers commented on comfort improvements. [...] Economic savings with the self-propelled platform ranged from \$128 to \$285 per acre for all tasks except tree training and pruning. [...] Orchard workers also commented [on] inherent savings, such as reduced fatigue.

While orchard platforms have been shown to increase efficiency, they have not seen widespread acceptance partly because they are designed for specific applications and are relatively expensive compared to the benefits accrued. Our goal is to demonstrate feasibility of technologies that will

improve efficiency in a number of applications by providing a common mobile platform that can be used in fully autonomous, partially autonomous, or fully manual modes. Consider, for example, an autonomous utility vehicle equipped with a boom of cameras and software that can locate, count, and size apples as it traverses an orchard. Such a vehicle could provide accurate crop load estimates weeks ahead of the harvesting season, thus increasing significantly the grower's capability to plan her harvesting labor pool. One variation would be an autonomous vehicle equipped with mechanical devices that perform gross, pre-thinning operations ahead of the finer, manual thinning; such a machine could decrease the total labor cost associated with the production process. Another variant is a vehicle that drives by itself at a creeping pace while farm workers standing on the vehicle can work on the canopy pruning and training the limbs at heights that would normally require ladders. Yet another variant is a vehicle that carries heavy crates of fruit as it is harvested, autonomously following workers in a row as well as autonomously traveling to the end of the row when the bins are full. Imagine now a single reconfigurable vehicle that could perform all the above tasks. We are developing a family of such vehicles, named Autonomous Prime Movers (Figure 1). They are described in Section III.



Figure 1. Autonomous Prime Mover, based on a Toro MDE eWorkman chassis. The two laser rangefinders on the front are used to find the tree rows closest to the vehicle. The information is passed on to the onboard computer, who finds the center of the row and guides the vehicle along it. Steering and wheel encoders complement the sensor suite and provide data for the computer to localize the vehicle along the row. We have implemented several vehicles of differing mechanical configuration, but since the sensing and computing architectures are the same, we are able to transfer autonomy software from one to the other with minimal rework.

II. RELATED WORK

Europe leads the development of vehicles specialized for specialty crops. These include orchard platforms (e.g. N. Blosi, Silverbull, Munckof) as well as small tractors (e.g. Goldini). We are aware of four orchard platform manufacturers in the US, including Blue Line in Washington and Phil Brown Welding in Michigan. Autonomous row following in orchard operations has traditionally been accomplished using computer vision. Billingsley and

Schoenfisch [4] demonstrated a vision guidance system for traversing cotton fields. Belforte et al. [3] designed a robotic system for navigating in greenhouses, and showed its effectiveness in spraying and precision application of fertilizer. Many groups in Japan use vision for a variety of agricultural applications, including rice planting, tilling, and spraying [14]. Laser rangefinders have been increasing in popularity in the robotics research community for sensing purposes. Numerous groups use lasers for obstacle detection and terrain modeling in off-road environments [12], [15]. Laser ranging is a mature technology and can provide robust sensing for the modes of operation proposed here [9], [10].

III. AUTONOMOUS PRIME MOVERS

Much of the rationale for our work is based on experience working with specialty crops as a part of the USDA-funded project "Comprehensive Automation for Specialty Crops" (CASC). One of the main CASC themes is "reconfigurable mobility" as personified by the development of an Autonomous Prime Mover (APM). The APM is an abstraction of an autonomous vehicle applicable to various scales of machines ranging from small ATVs to large orchard platforms and tractors. The APM is analogous to the personal computer. Personal computers vary in terms of computing power and size, but share a common infrastructure (processors, hard disks, memory, operating system, etc.) that allows users to exchange applications and documents. The idea with the APM is the same: irrespective of the vehicle size, features, or carrying capacity, applications developed for one should be seamlessly transferable to others (as long as the appropriate hardware is available). The current APM is already capable of autonomous driving along a row of trees, turning at the end of the row and entering the next one. Row following is conducted at the center of the row independently of the row width (for sensing or mowing) or at a pre-defined distance from the trunk line (for pruning, thinning, or spraying) [9], [13]. Experience with deployment of the APM in commercial orchards in Pennsylvania and Washington has resulted in understanding critical needs that a successful platform must meet in addition to being multifunctional and low-cost:

- the vehicle must fit into rows of modern high-density orchards without damaging fruit. Ideally, the vehicle would be at most 48" wide at 3 ft. above the ground, but would have a platform that can be widened as much as 8 ft. to deal with wider plantings;
- smooth forward speeds from 0.2 to 5 mph for operation within crops as well as a higher speed mode of 10-15 mph while being driven to blocks;
- stable on cross slopes of up to 10°;
- full workday endurance before refueling/recharging;
- operation in temperatures varying from 10 to 100° F;
- electric-battery powered for smooth/quiet operation as well as reduced emissions;
- self-steered and easily operated by farm workers.

Working with growers and extension educators, we came up with four modes of operation, each corresponding to distinct classes of tasks:

Transit mode: In this mode, the vehicle is manually controlled for transport from a storage location to a place where it can be deployed. It should be possible for the vehicle to be controlled as easily as a golf cart or a utility vehicle at speeds of up to 15 mph.

Scaffold mode: In this mode, farm workers stand on the platform and perform tasks like pruning, thinning, and tree training and tying, moving at speeds of 0.2 mph. This is particularly relevant in tall spindle canopies such as in high-density apple orchards. With the attachment of a dry bin filler, such a platform could also be used during harvest. The standing locations on the vehicle have to be designed such that the average worker isn't often required to raise his/her arm above shoulder level. Ideally, in this mode the vehicle is self-steered, eliminating the need for a driver. Since a farm worker will be onboard, this mode can be accomplished with simple sensors and can leave the uncommon cases (e.g., an apple bin or a stopped vehicle on the way) for the human worker to resolve. Minor adjustments like distance to the canopy can be made easily via a joystick.

Mule mode: In this mode, farm workers walk along the row, tending and harvesting and placing fruit in bins on the vehicle. When bins are filled, the vehicle can be directed to autonomously transport them to the end of the row, moving at speeds of 2-3 mph. Such operation is relevant, e.g., in table grapes where loaded crates must be carried back to the end of the row where other workers load the crates onto a truck. This mode will require a more sophisticated interface between the workers and the machine because the workers will not necessarily be on or near the machine. Also, the vehicle will have to travel some distance along the row on its own and then back to where the workers are picking, stopping automatically when either the end of the row is reached or it has approached close enough to a worker.

Pace mode: In this mode, the vehicle autonomously drives an entire block at a time without requiring any further interaction. No operator will be on the vehicle or will be necessary to interact with it until the assigned task is complete. This mode is useful in tasks such as spraying (both block and selective), mowing the cover crop in between the rows and inspecting the canopy for disease and pests. This mode could also be used to collect data for yield estimation because this requires an accurate control of statistical sampling—for example, the vehicle might be required to move fast for 50 yards and then slow down to take detailed data (typically with a camera tuned to specific wavelengths) for 5 yards at 0.5 mph, and then repeat this pattern over the block. In this mode, a camera, sprayer and mower could be easily swappable both physically and in software much like a printer, scanner, or mouse is attached to a personal computer.

IV. VEHICLE DEVELOPMENT

We believe that the successful commercial deployment of APM-like vehicles will be the result of a progression of steps aimed at proving the system's feasibility, incorporating robustness to the target environment and users, and making them easy to use and maintain—in addition to ensuring they are cost effective. In developing the APM concept we started with a proof-of-concept vehicle that could drive in “easy” environments with mostly flat terrain and wide rows; added robustness to terrain and canopy variation; and finally created interfaces for the vehicles to be operated by growers, without engineering support.

A. Sensing Suite

The first APM prototype, named Laurel, was equipped with laser sensors, driving and steering wheel encoders, onboard computers, and a software suite to enable the vehicle to drive between rows of trees and turn around at the end of rows [9], [13]. Laurel's original design had two laser rangefinders located on the corners of the vehicle about one foot off the ground (Figure 1). The lasers scan 180° in a horizontal plane and provide distance to objects in front of them. We oriented the lasers to give a 270° field-of-view in front of and to the sides of the vehicle.



Figure 2. (Top) Typical orchard map produced by the planar laser rangefinders. The vehicle is represented by a gray rectangle, with the red part indicating the front. The white dots are the rows of trees or other objects within the lasers' field-of-view (weeds, terrain, etc.). The yellow and blue lines represent the best reconstruction of the tree rows given the laser data; note how slopped terrain to the right of the vehicle introduces errors in the tree row estimation, which in turn leads to inconsistent row following. (Bottom) Laurel's new configuration. The laser rangefinder mounted on the center of the front grille is the only one used for row following.

This is a simple, low-cost setup that allows the autonomous system to see the world around it, though only the objects that are one foot high relative to the vehicle. This includes not only tree trunks and canopy, but weeds and sometimes the ground itself when the terrain slopes upwards (Figure 2, top). To the autonomous system such obstacles are indistinguishable from trees and make row detection difficult.

To overcome this limitation, we modified Laurel to use only one laser rangefinder (270° field-of-view) for row following.. We mounted the sensor centered left-right and at a height just above the vehicle's hood (Figure 2, bottom). This configuration has two advantages. First, because the laser is above the hood we can take advantage of its increased field-of-view. Second, the sensor at this height is less susceptible to the spurious obstacles mentioned above.

B. Row Following

Our first row detection system used a Hough transform [6], which looks for the most likely pair of parallel lines in the laser range data. In normal circumstances, when the laser only sees the tree trunks and canopy, the two lines of tree canopy are the most likely. The autonomous system then finds the midway line between the two, and that is selected as the desired path to drive. When, however, spurious obstacles are present, the system may perceive them as being part of a good pair of parallel lines, which results in the vehicle drifting off center, sometimes into the canopy (Figure 3, top). Although the new laser height mitigated this problem by presenting fewer such obstacles, we required a more robust row detection system.

We focused then on developing a row detector that uses a particle filter [5], [8]. The particle filter makes multiple guesses of where the tree rows could be, and scores each guess by how much it agrees with the laser data. Furthermore, high-scoring row lines are kept from one iteration of the detection to the next, so that detections get better over time. When spurious obstacles appear, the filter remembers the previous row lines that had been detected, and can select the correct row (Figure 3, bottom).

C. Row Turn

We designed turn methods to make row entry smooth and reliable. We implemented a path planner that takes the nominal starting point of the vehicle in the next row and generates a smooth path that respects the vehicle's steering constraints [7]. The planner optimizes the path to avoid hard steering angles that stress the motor, while minimizing the total path distance (Figure 4, top left). This planner aligns the vehicle with the row as it enters, so the vehicle is pointed along the row early. Before the vehicle completes the turn the autonomous system runs a row detection to refine its target entry point (Figure 4, top right).

Frequently, in order for the vehicle to be aligned with the row before entering, the planner creates a large-radius, bulb-shaped turn. This turn is not always possible due to space constraints. Therefore we have also developed a three-point,

"K" turn. The vehicle first turns past the row, then backs up and points towards the row, and finally enters the row (Figure 4, bottom). During the third leg of the turn the system runs row detection to refine the final goal point.

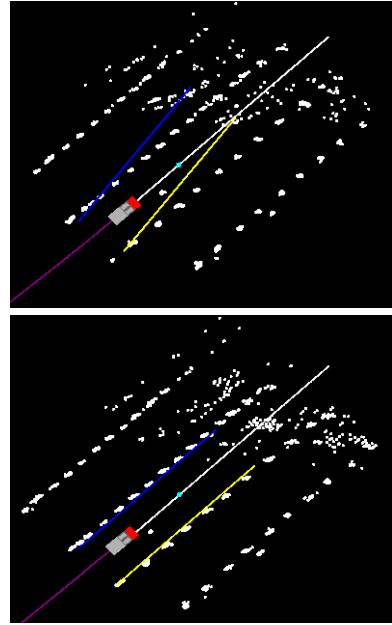


Figure 3. Comparison of the performance of two row detection methods among tall weeds or uneven terrain. Since we use a planar laser, the returns can come from the terrain or weeds rather than from the trees. (Top) The Hough transform method used initially is confused by the unexpected data, resulting in bad row detections. (Bottom) The new particle filter method correctly ignores the additional data and finds the tree rows.

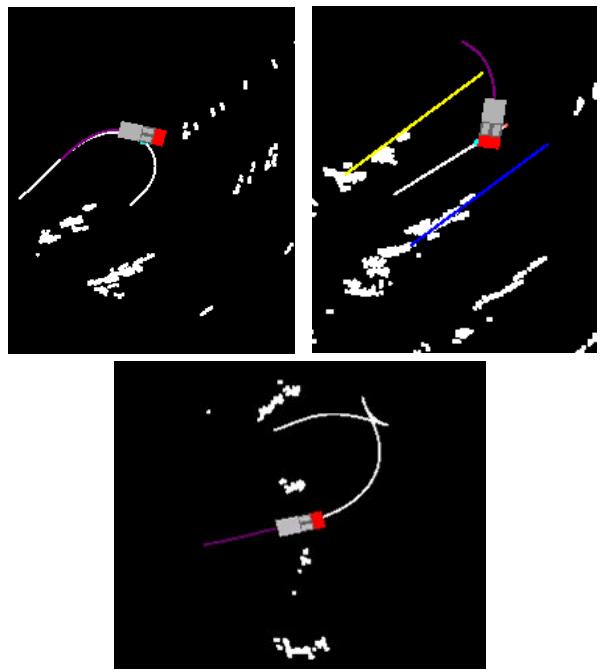


Figure 4. (Top left) The row turn planner generates a smooth trajectory to follow, staying away from hard steering angles that stress the motor, while aligning the vehicle to be able to see into the row before entry. (Top right) Once the vehicle has turned towards the row, the system performs a row detection to refine the desired path of the vehicle. (Bottom) This row is too narrow to smoothly turn into. The system plans a "K"-turn, where the vehicle turns past the row, backs up, and then enters. In the third leg of the turn the system performs a row detection to refine the final desired path.

We successfully tested both of the new turn methods. The limitation at this point is that we need to manually program which type of turn to do, smooth or three-point. We also tell the system how much space it will have in which to turn the vehicle (for planning collision-free paths). In future work we will program the system to autonomously assess the drivable space and decide which turn maneuver to execute.

D. Row Entry

When the vehicle reaches the end of one row it must turn into the next row. Due to inconsistencies in tree plantings and in the vehicle's location estimate, the system does not rely solely on the prescribed row width for its target position. Instead, as it is turning, it runs the row detection to find the lane to go into. Our initial turning method made a sharp turn towards the next row, and ran row detection when the vehicle was perpendicular to the row. From this point, however, the canopy of the nearest trees can block the far row, giving the autonomous system little data to work with (Figure 5). Detections from this point were unreliable, with a success rate less than 50% in dense canopy.

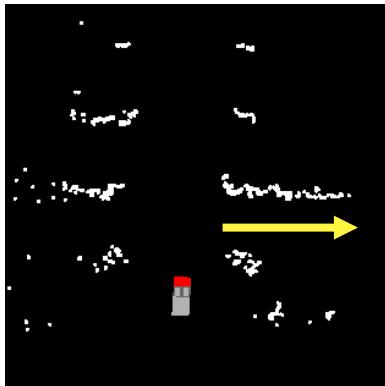


Figure 5. Laurel attempts to enter the row lane indicated by the yellow arrow. The canopy of the trees at the near row, however, blocks the rest of the near edge and most of the far edge. The row detector has little data to work with in this situation.

E. Usability Aspects

Having completed the development of the base autonomy software, we started looking toward commercial deployment of the APM vehicle family. The first step was to create vehicles that can not only survive the harsh orchard environment year round, but also be operable by growers and farm workers. To that end we developed three new vehicles. One is used in *pace* mode. The other two are two-person *scaffold* mode platform vehicles meant to be operated by workers standing on the vehicle performing tasks that would be normally conducted by repeatedly going up and down a ladder (Figure 6).

The computing demands of our system could be satisfied by an average personal computer. Office computers, however, are not robust to the high temperatures found in all-day orchard operation during the summer, nor are they sealed to prevent damage from dust or rain. To operate a vehicle outdoors continuously we require a specialized

computing solution. We chose an embedded computer that is completely sealed and rated to a temperature of 105° F. The embedded computer runs the autonomy software, receiving data from the laser rangefinder and sending commands to the vehicle's drive-by-wire controller.



Figure 6. One use for self-driving vehicles in orchards is in “scaffold” mode in which farm workers work from a platform performing tasks such as pruning and thinning while the vehicle drives by itself at low-speeds.

The user controls the system from an interface (Figure 7) that communicates with the embedded computer wirelessly. The interface has been specifically designed so that it can be operated by workers with minimal training using intuitive controls. From this interface, the operators can control vehicle speed, offset from the center of the row, and select between continuous drive, where the vehicle runs at constant speed until stopped by the user, or intermittent drive, where the vehicle goes a certain distance and stops on its own. The interface also allows for manual control via a joystick.



Figure 7. Interface with controls for vehicle speed and offset from the center of the row, as well as a joystick for manual driving and a message panel for feedback.

Finally, we needed to consider remote maintenance of the vehicles. As we continue to improve the software we need a way to provide updates to the users. Also, in case of abnormal vehicle behavior we need a way to access data logs on the embedded computer. To solve these problems we

integrated a cellular modem in each vehicle. The embedded computer runs a script such that, when the modem is turned on, the script automatically creates a connection to a server at Carnegie Mellon. This way, we can provide updates and perform remote debugging seamlessly to the user.

V. FIELD TESTS

We conducted tests of the APM autonomous navigation system in research and commercial orchards in Pennsylvania and Washington states. To date we have logged over 330 km, including over 100 km in fully developed orchards to test system's ability to handle different blocks and canopy types. The vehicles successfully drove several different blocks, despite differences in row width, canopy training, and growing systems, using as the only *a priori* information the number of rows and length and width (Figure 8)

In particular, the two platform-equipped vehicles drove 30 km autonomously in commercial orchards carrying workers conducting a variety of operations. Preliminary results indicate efficiency improvements up to 58% compared to workers on ladders (Figure 9). Additionally, the preliminary results indicate energy savings of 97%, with the electric APMs coming in at \$0.16/acre versus diesel platforms at \$5.95/acre.



Figure 8. A sampling of the blocks driven by Laurel in 2010. Growing systems varied from recent fruiting wall plantings to completely standalone, untrained trees. All pictures are from blocks at the Penn State Fruit Research and Education Center

VI. CONCLUSION

As fruit production moves from large, tall trees to high-density fruit walls, the potential for automation increases

significantly. While we may be many years away from robotic fruit picking that is economically viable (see, e.g., [1]), we already have the technology necessary to create cost-effective autonomous orchard vehicles that can augment or automate tree pruning and training, blossom and fruit thinning, fruit harvesting, mowing, spraying, sensing, etc. In fact, orchard platforms are a daily reality in many tree fruit production operations in Europe, and their presence in the U.S. has been growing steadily. Here we have described a family of autonomous orchard vehicles and platforms that can be reconfigured and retasked year-round for a variety of operations. To achieve the ultimate goal of turning these vehicles into commercial products, we follow a stepped approach in which feasibility is followed by robustness and usability. As of the end of 2011, we had completed over 330 km of autonomous orchard traversals in experimental and commercial orchards in Pennsylvania and Washington.

VII. REFERENCES

- [1] Baeten, J., Donne, K., Boedrij, S., Beckers, W. and Claesen, E. "Autonomous fruit picking machine: A robotic apple harvester," Intl. Conf. on Field and Service Robotics, France, 2007.
- [2] Baugher, T., Schupp, J., Lesser, K., Harsh, R.M., Seavert, C., Lewis, K. and Auvil, T. "Mobile Platforms increase Orchard Management Efficiency and Profitability," Acta Hort. (ISHS) 824:361-364, 2009.
- [3] Belforte G., R. Deboli, P. Gay, and P. Piccarolo. "Robot design and testing for greenhouse applications." Biosystems Engineering. 95(3): 309-321, 2006.
- [4] Billingsley J. and M. Schoenfish. "The successful development of a vision guidance system for agriculture." Computers and Electronics in Agriculture. 16(2): 147-163, 1997.
- [5] Dellaert, F., D. Fox, W. Burgard, and S. Thrun. "Monte Carlo localization for mobile robots." IEEE International Conference on Robotics and Automation, 1999, pp. 1322-1328.
- [6] Duda, R. O. and P. E. Hart. "Use of the Hough Transformation to Detect Lines and Curves in Pictures." Communications of the ACM, 15: 11-15, 1972.
- [7] Ferguson, D. Path Generation and Control for End of Row Turning in an Orchard Environment. Master Thesis, Carnegie Mellon University Robotics Institute, 2010.
- [8] Gustafsson, F., F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. "Particle filters for positioning, navigation, and tracking." IEEE Transactions on Signal Processing, 50(2):425-437, 2002.
- [9] Hamner, B., S. Singh, M. Bergerman. "Improving orchard efficiency with autonomous utility vehicles." American Society of Agricultural and Biological Engineering Annual Convention Proceedings, 2010.
- [10] Libby, J. and Kantor, G. "Deployment of a point and line feature localization system for an outdoor agriculture vehicle." IEEE International Conference on Robotics and Automation, pp. 1565-1570, May 2011.
- [11] Oetomo, D., Billingsley, J. and Reid, J. F., "Editorial: Agricultural robotics." Journal of Field Robotics, 26: 501-503, 2009.
- [12] Peynot T., J. Underwood, and S. Scheding. "Towards reliable perception for unmanned ground vehicles in challenging conditions." IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1170-1176, 2009.
- [13] Singh, S., M. Bergerman, J. Cannons, B. Grocholsky, B. Hamner, G. Holquin, L. Hull, V. Jones, G. Kantor, H. Koselka, G. Li, J. Owen, J. Park, W. Shi, J. Teza. "Comprehensive Automation for Specialty Crops: Year 1 Results and Lessons Learned." Journal of Intelligent Service Robotics, Special Issue on Agricultural Robotics, 3(4): 245-262, 2010.
- [14] Torii, T. "Research in autonomous agriculture vehicles in Japan." Computers and Electronics in Agriculture. 25(1-2): 133-153, 2000.
- [15] Wellington C., A. Courville, and A. Stentz. "A Generative Model of Terrain for Autonomous Navigation in Vegetation." International Journal of Robotics Research. 25(12): 1287-1304, 2006.

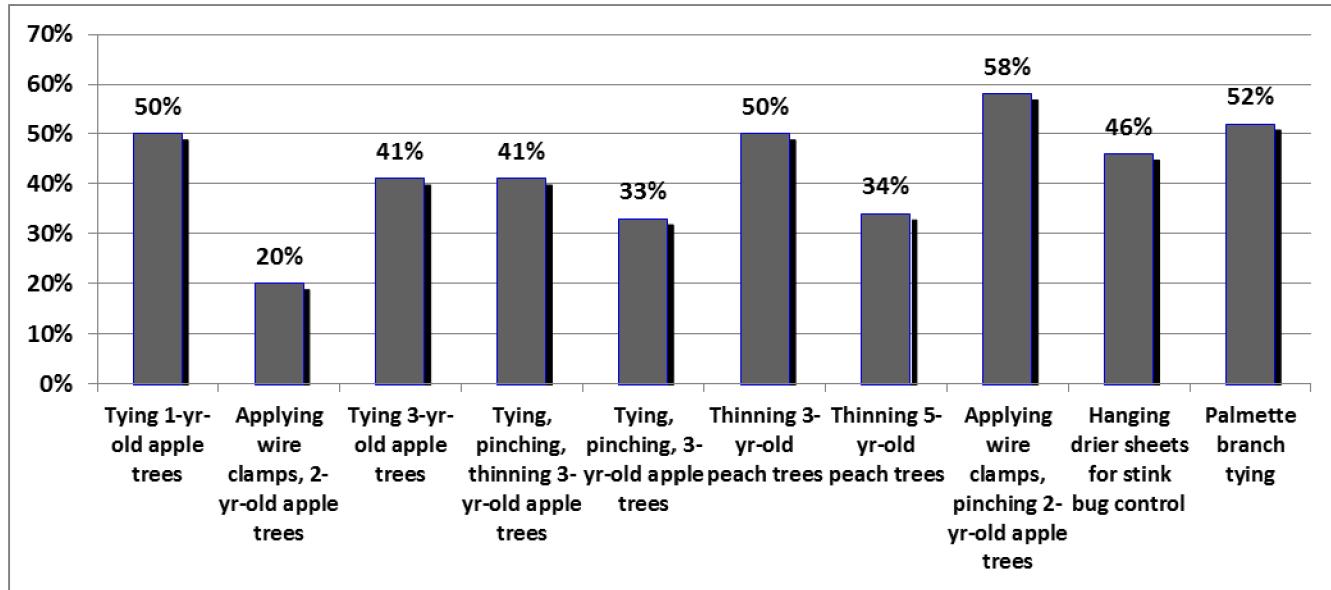


Figure 9. Increase in efficiency for a variety of tasks performed using the vehicle shown in Figure 6 over ladders, in trials conducted at the Penn State Fruit Research and Extension Center. Videos of workers on the platform and on ladders working side by side, and workers using the vehicle during harvest as a “bin dog.” can be found at <http://bit.ly/wUxf3b>.