

Devices, Systems, and Methods for Automated Monitoring enabling Precision Agriculture

Jnaneshwar Das*, Gareth Cross*, Chao Qu*, Anurag Makineni, Pratap Tokekar, Yash Mulgaonkar, Vijay Kumar

(*the first three authors contributed equally to this manuscript)

GRASP Laboratory, University of Pennsylvania
{*djnan,gcross,quchao,makineni,tokekar,yashm,kumar*}@seas.upenn.edu

Abstract—Addressing the challenges of feeding the burgeoning world population with limited resources requires innovation in sustainable, efficient farming. The practice of precision agriculture offers many benefits towards addressing these challenges, such as improved yield and efficient use of such resources as water, fertilizer and pesticides. We describe the design and development of a light-weight, multi-spectral 3D imaging device that can be used for automated monitoring in precision agriculture. The sensor suite consists of a laser range scanner, multi-spectral cameras, a thermal imaging camera, and navigational sensors. We present techniques to extract four key data products—plant morphology, canopy volume, leaf area index, and fruit counts—using the sensor suite. We demonstrate its use with two systems: multi-rotor micro aerial vehicles and on a human-carried, shoulder-mounted harness. We show results of field experiments conducted in collaboration with growers and agronomists in vineyards, apple orchards and orange groves.

I. INTRODUCTION

The world population is estimated to reach 9 billion by 2050. Sustainable food production remains a global concern. The difficulty of meeting global nutritional needs is further compounded by endemic crop diseases and scarcity of water for irrigation. As a consequence, efficient resource management and persistent and timely monitoring of crop health are becoming increasingly crucial. In this paper, we present the design of a sensing device that allows for effective and timely monitoring of crop in farms. This device can be used in hand-held mode as well as mounted on low-flying, aerial robots, and allows us to extract data products that are crucial for efficient and sustainable farming.

Traditionally, remote sensing satellites and airborne sensing with winged aircraft have allowed scientists to map large farmlands and forests through acquisition of multi-spectral imagery and 3-D structural data [1], [2]. However, data from these platforms lack the spatio-temporal resolution necessary for precision agriculture. For example, a typical remote sensing satellite image may have a pixel resolution of hundreds of meters, and airborne sensing may provide resolution of a few meters. Monitoring orchard or vineyard health however requires data at a centimeter scale – the resolution necessary for observing stems, leaves, and fruits. Unmanned Ground Vehicles (UGVs) have been the first step towards close-range monitoring of high-value crops. They

can carry a variety of bulky sensors such as LiDAR for volumetric mapping [3], and Ground Penetrating Radar (GPR) and electrical conductance sensors for precise soil mapping. However, due to the mobility constraints of unstructured farms, it is infeasible to use UGVs for rapid and persistent monitoring. Additionally, ground vehicles are intrusive. We assert that the dense mapping of orchards and vineyards is a task best served by aerial platforms and hand-held sensors.




Fig. 1: The sensor suite was designed to acquire multi-modal, multi-scale data, while mounted on an Unmanned Aerial Vehicle (UAV), a steadicam, or through handheld human operation. The sensor suite consists of multi-spectral cameras, a thermal imaging camera, a laser range scanner (LiDAR), and navigational sensors, providing a rich spatially registered dataset for extraction of a wide range of plant properties of interest.

A significant portion of prior work in robotics for agriculture has tackled the issues of farm preparation, crop harvesting, and weeding [4], [5], [6], [7]. In recent years, there has been a growing interest in the use of imaging sensors for monitoring plant physiology and morphology. LiDAR scanners mounted aboard ground vehicles and tractors have been used to extract morphological properties such as canopy volume and leaf area [8]. Thermal and multi-spectral imagery in the red and near-infrared bands have been used to monitor plant health, guiding pruning management and

fertilization [9], [10]. Thermal and multi-spectral cameras have been used together onboard UAVs to estimate a range of plant properties related to photosynthetic efficiency and water stress [11]. However, the size, weight, and costs of existing multi-modal imaging systems inhibit large-scale deployment onboard UAVs. It is hence desirable to develop a portable, low-cost, compact, and light-weight imaging system along with agile deployment methodologies to help growers observe farms effectively. Additionally, current work in precision agriculture does not sufficiently address the development and evaluation of a comprehensive data analysis and visualization framework that can provide growers with actionable intelligence, and help them interpret the resulting data products efficiently.

To this end, in this work we present a robotic sensing methodology that exploits a light-weight and portable multi-spectral 3-D imaging system (Figure 1) that can be mounted onboard Unmanned Aerial Vehicles (UAVs) for rapid sensing in unstructured farmlands (Figure 2a). The sensor suite can also be deployed in a harnessed mode (Figure 2b), providing farm owners with the option of aerial deployment or manual deployment depending on the structure of the environment. For example, many close-range sensing requirements can only be met in the harnessed mode, whereas aerial deployment can provide rapid synoptic imagery from the top. The sensing modalities of the sensor suite were selected to monitor a range of plant physiological and morphological properties, and we have developed approaches to extract actionable intelligence from the data acquired by the system. Specifically, we present methods to estimate four properties that are of interest to farm owners – plant morphology, plant vigor, leaf area, and fruit counts. Results from multiple field trials are discussed in the paper to demonstrate the deployment modes and data products, highlighting the applicability of our system for a variety of precision farming tasks.

The key contributions of our work are as follows. First, we combine different sensing modalities on a self-contained, lightweight, and compact platform. The closest system to our work is equipped with onboard LiDARs, a spectrometer, and a thermal camera, developed for measuring tree characteristics [12]. However, in addition to possessing the sensing modalities carried by this sensor, our system can be mounted on multi-rotor UAVs with payload constraints, as well carried by human operators for handheld use. Second, we use machine learning techniques to estimate properties of interest such leaf area, and fruit count from the data acquired by the sensor suite. Finally, we use state of the art tools to visualize data products generated by our system.

The rest of the paper is organized as follows. In Section II, a detailed description of the hardware design and software stack is provided. In Section III we describe techniques to extract four data products from the sensor suite and presents results from multiple field trials. We conclude in Section IV along with a discussion of future research directions.

II. SYSTEM DESCRIPTION

In this section, we describe the hardware design, deployment modes, and software stack of our system.



(a) The sensor suite mounted on a UAV, scanning plant canopies from the top with the sensors pointing downwards.



(b) The sensor suite mounted on a hand-carried steadicam, providing side view of plants (towards the left in this image, pointing at a row of apple trees). The white dashed lines show the vertical field of view of the system. This arrangement provides rich data on fruits, canopy, and trunks. Such data can also be acquired with the sensor suite mounted on a UAV flying between trees.

Fig. 2: Our system is currently being deployed in two modes — mounted on a hand-carried steadicam, and on a UAV.

A. Hardware design

The sensor suite is shown in Figure 1. It consists of an array of science and navigation sensors, an onboard computer, a wireless communication link, and onboard batteries. A Hokuyo UST-20LX laser scanner provides high-resolution laser scans for monitoring plant morphology. Two monochrome Matrix Vision BlueFox cameras equipped with narrow-pass filters at 670nm and 800nm allow for calculation of the Normalized Difference Vegetation Index (NDVI) — an important indicator of plant vigor. A FLIR A35 thermal

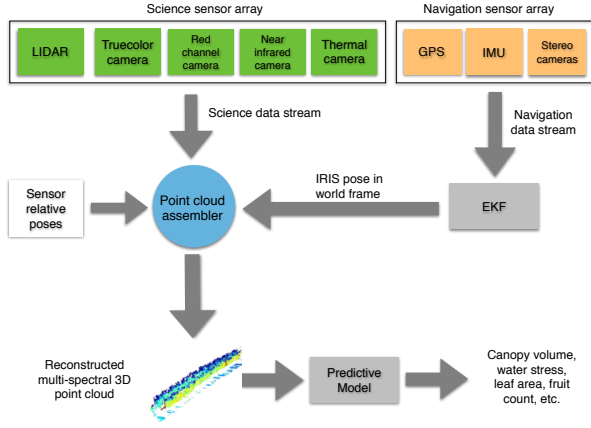


Fig. 3: The data processing pipeline starts with the sensor streams from the science and navigation array, which are fused to generate a multi-spectral 3-D point cloud.

imaging camera collects temperature readings at a resolution of 320x256. Finally, an RGB BlueFox camera is used to acquire true-color data, which is used for fruit-counting and visual inspection. The navigational sensor array consists of a Microstrain 3DM-GX4-25 Inertial Measurement Unit (IMU), a μ Blox precise point positioning (PPP) GPS sensor, and two single-channel BlueFox cameras for stereo visual odometry. An onboard Intel core i5 computer logs data from all the sensors and facilitates communication with the vehicle it is mounted on, as well as with a ground station through a wifi link.

The sensor package has a total weight of ~ 1.6 kg, and dimensions of 40 cm x 13 cm x 13 cm, making it lighter and more compact than sensor suites with similar capabilities. A carbon fiber frame supports a poly-carbonate base plate on which all sensors are mounted. Power is delivered by two 2700 mAh lithium polymer batteries. We have observed an endurance of one hour during the deployments. The sensor arrangement has been improved over multiple field trials for reliable operation in different deployment modes.

B. Deployment

The sensor suite has been designed to be versatile, capable of being mounted on UAVs or used directly in a hand-held mode. In our work, we have tested the sensor suite on a DJI Innovations S800 multi-rotor UAV, and a steadicam rig (Glide Gear DNA 5050 Vest And Arm Stabilization System Pro). Figure 2a shows the sensor suite mounted on the UAV while scanning orange tree canopies from above. Figure 2b demonstrates the sensor suite being used with the steadicam rig in order to scan apple trees looking from the side while the operator walks at a normal pace. Mechanical stabilization of the steadicam improves the quality of the recorded data, resulting in higher precision of 3-D reconstruction.

C. Software stack

The data acquired by the sensor suite is processed in multiple stages. The entire pipeline is illustrated in Figure 3. The Robot Operating System (ROS) [13] forms the backbone of the data processing pipeline, facilitating sensor data logging and sharing of information between various processing nodes. The data from the navigation sensors are used by an Extended Kalman Filter (EKF) to generate pose estimates for every sensor on the platform. We employ the *error-state* formulation of the EKF, similar to the form documented in [14]. A point cloud assembler uses the EKF pose estimates and the known relative poses between sensors to reconstruct a multi-channel 3-D point cloud. The point clouds are also converted to an octree [15] representation for efficient storage and analysis.

D. Calibration

Since the sensor suite is a platform for performing detailed and highly accurate environmental reconstruction, it is important that both the state estimation and scientific sensors be properly calibrated. Prior to each deployment, we perform a complete system calibration, wherein the stereo and sensing cameras are calibrated relative to the IMU. We employ the toolbox from [16], [17] to perform both camera-system calibration (stereo, color and multi-spectral) and spatial/temporal calibration of camera and IMU.




Fig. 4: Camera system calibration between thermal camera and left stereo camera. The circle grid pattern has been heated before calibration

In addition, we calibrate the intrinsic and extrinsic parameters of thermal camera relative to the stereo pair. This procedure is complicated by the fact that the thermal camera cannot observe visible wavelengths, and produces only a low-resolution image (320x256). Thus a standard chessboard pattern cannot be used to calibrate a thermal camera directly. In order to address this problem, we leverage a strategy similar to [18]. An ordinary circle grid pattern printed on paper is illuminated by a hot lamp, producing a pattern which is discernible in both the long-wave IR and optical regions of the spectrum, as shown in Figure 4. This approach allows us to calibrate thermal camera itself as well as with other cameras without requiring any complicated calibration device.

III. DATA PRODUCTS AND FIELD TRIALS

To highlight the applicability to farm management, we discuss four example data products obtained during the field trials using our system. These are: reconstructing plant

morphology, computation of plant vigor, estimation of leaf area, and automated fruit counting or yield estimation using remotely sensed data. Targeted at improving farm management, these applications demand extraction of actionable intelligence from the data collected by the sensor suite, mounted either on a UAV or on the harness. Plant vigor, measured through the normalized difference vegetative index (NDVI) facilitates decision-making for fertilization. Accurate estimation of leaf area has the potential to improve pruning and spraying management. The capability to estimate yield accurately will enable growers to plan labor for harvesting, and storage for harvested fruits. Both applications require training of predictive models that use the acquired data to estimate a property of interest, i.e., leaf area or fruit count.

In this section, we first discuss the characteristics of reconstructed plant data, followed by description of the other three data products extracted by the software stack. We demonstrate the utility of our system through data acquired and processed during multiple field trials. We used the sensor suite onboard a UAV and on a shoulder-mounted harness in a vineyard in Galt, California, an apple orchard in Biglerville, Pennsylvania, and at an orange grove in Orange Cove, California.

A. Plant Morphology


Figure 5 shows a reconstructed point cloud of a row of grape trees from the vineyard in Galt, California. In Figure 5b, features of the environment, e.g., the canopy, trunk, ground have been highlighted.

In the apple orchard at Biglerville, PA, we scanned a row of apple trees of different sizes across two field trials. A 3-D reconstruction from the first trial was carried out to scan semi-dwarf apple trees. The reconstruction and canopy dimensions are shown in Figure 6. The second trial was carried out to scan smaller, dwarf apple trees. Thermal camera data was projected on this point cloud as shown in Figure 7. Warmer sections are showed in red, and cooler sections in blue.

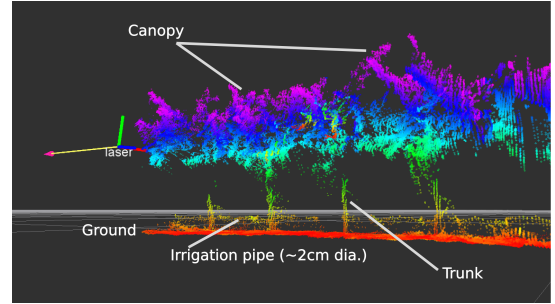
The point clouds from both trials were analyzed to determine canopy volume characteristics. Figure 8 shows two dwarf apple trees from the row of trees scanned. A comparison with Figure 6b shows the difference in dimensions between semi-dwarf and dwarf apple trees – the semi-dwarf trees have a height of ~ 3.5 m, and the dwarf trees a height of ~ 2.5 m.

B. Plant vigor

Normalized Difference Vegetative Index (NDVI), an indicator of plant vigor, was computed using multi-spectral imagery acquired by the sensor suite onboard the UAV. For pixel data corresponding to a multi-spectral image, $NDVI = (NIR - VIS) / (NIR + VIS)$, where $NIR = 800\text{nm}$ (i.e., near-infrared) and $VIS = 670\text{ nm}$ (i.e., visible). Figure 9 shows a visualization of NDVI map generated using multi-spectral data acquired by the UAV. These maps will enable growers to plan fertilization and mitigation in response to stresses observed in the NDVI imagery.



(a) Reconstruction of a section of the row of grape trees spanning 30 m. The color shows the intensity of laser scan echo, red being higher intensity.




(b) Various features of the reconstructed trees. Color shows the height, red being ground.

Fig. 5: Reconstructed 3-D point cloud of a row of grow trees at the vineyard in Galt, CA. Data was collected with the sensor suite mounted on the steadicam facing the side view of the grape trees.


C. Leaf area estimation

Estimation of the total leaf area of a plant is important for management of fertilization, pruning, and spraying. We use the data from the sensor suite to obtain accurate leaf area for target trees. To develop and test our leaf area estimation methodology, we defoliated two dwarf apple trees in the apple orchard in Biglerville, PA, and concurrently carried out scans with the IRIS imaging system mounted on a steadicam. The defoliated leaves were used to determine ground truth leaf area, necessary for developing the algorithm to estimate leaf area from remotely sensed data. Complete tree defoliation was done in four passes, with leaves from two octants (illustrated in Figure 10) removed after each pass. This emulates four canopy types for each tree – with all leaves, with 1/4 leaves removed, with 1/2 leaves removed, and with 3/4 leaves removed. The defoliated leaves were collected and brought back to the lab, and image processing algorithms were used to compute ground truth leaf area for each octant of the defoliated trees.

To carry out this estimation, the leaves were flattened on a table using a polycarbonate sheet and photographed with a DSLR camera. A calibration target placed at the corner of the table provided scale and allowed estimation of camera pose to carry out perspective correction prior to the use of a segmentation algorithm. In total, sixteen ground truth leaf area estimates were obtained, one from each octant of




(a) A section of five semi-dwarf apple trees (annotated as section 1), reconstructed using the sensor suite.



(b) Profile view of a semi-dwarf apple tree at the edge of a row of trees. The 3-D point cloud was reconstructed through the sensor suite, and illustrates the tree dimensions.


Fig. 6: Reconstructed 3-D point cloud of a row of semi-dwarf apple trees in an orchard in Biglerville, PA.



(a) Image of leaves placed on a table, captured through a DSLR camera.



(b) Unwarped image after a perspective transform using the camera calibration board placed in the right top corner of the table.



(c) Individual leaves detected using an image segmentation algorithm, leading to the computation of individual and total leaf area.

Fig. 11: The data processing sequence for extracting ground truth leaf area for each octant of defoliated leaf.




Fig. 7: Thermal data projected onto the 3-D point cloud of a row of dwarf apple trees in Biglerville, PA. The color corresponds to temperature, red being hot. The right side of the tree row was facing west, a few hours before sunset. Hence, it is warmer than the other parts of the canopy.

each tree. Four IRIS sensor passes carried out preceding each stage of defoliation resulted in eight data points that were used to evaluate the leaf area estimation algorithms. The point cloud corresponding to each defoliation stage for each tree was converted into an occupancy grid with 5cm resolution, taking into account the location of the LiDAR on the sensor suite, and the resulting paths of the laser beams. Figure 12 shows the occupancy grids for the two trees. k-fold cross validation on a linear model of various choices of input features resulted in the best leaf-area estimates




Fig. 8: Two trees extracted from the dwarf apple tree point cloud to illustrate the tree geometry.

using weighted occupied-voxel density. Here, an occupied voxel has an occupancy probability greater than 0.5, and volume of the tree was taken to be the bounding box of all occupied voxels. Weighting of the voxels was done using the occupancy probability. We call this metric the LiDAR area index. Figure 13 shows the correlation between the LiDAR area index and the true leaf area for each of the data points. This choice of input feature (i.e. LiDAR area index) provided the best performance, with an R-squared value of 0.82. The result demonstrates the use of the IRIS sensor suite in estimation of leaf area for rows of trees.

Although evaluation was carried out in the harnessed



Fig. 9: ROS Rviz screenshot showing the UAV in flight above a row of orange trees in Booth Ranches, CA. The panels on the right show RGB, NDVI, and IR images acquired by the IRIS sensor suite onboard the UAV (top to bottom). NDVI imagery is generated using data from the multi-spectral cameras of the sensor suite.




Fig. 10: Each octant of the two dwarf apple trees was defoliated separately for ground truth leaf area measurement.




Fig. 12: Occupied voxels (probability > 0.5) of the occupancy grids generated from the point clouds.




Fig. 13: Correlation between LiDAR area index and measured true leaf area for the data points corresponding to the emulated canopy stages of the two trees. An R-squared value of 0.82 was observed, suggesting good predictability.

mode, the methodology can be extended to rapid scanning with IRIS mounted on UAVs. Our leaf area estimation methodology can be used to rapidly estimate the leaf area of trees in a farms, enabling precise fertilization, spraying, and pruning.


D. Automated fruit counting

Accurate automated fruit counting will enable growers to determine storage and labor needs in prior to harvest. Our work uses data acquired from the sensor suite to generate fruit count to provide this capability. Here, we discuss the methodology and results from an experiment at the citrus grove in Orange Cove, CA. We scanned 26 rows of trees of a block of size 180m x 180m from both sides of each row using the harnessed IRIS sensor suite. An algorithm was developed to generate fruit counts for the rows using the data acquired by the sensor suite. Although discussed in the context of counting oranges, our approach can be extended to counting of other fruits such as apples, peaches, as well as clustered fruits such as grapes and blueberries.

Our fruit counting approach consists of two steps, fruit detection followed by fruit tracking. We carry out fruit detection using a support vector machine (SVM) classifier that uses different color spaces to classify each pixel in an image as originating from a fruit or not. A series of images were labeled to annotate regions that have a fruit enclosed (an orange in our use case). The training dataset of images were used to train the SVM classifier with candidate pixel colorspace values as input. For estimating fruit yield from farms, it is necessary to generate a running count as the sensor suite carries out scans of trees. To do so, we have developed a fruit tracking algorithm that tracks fruits detected in a stream of images. To keep track of fruits that have already been detected in previous frames, we compute the optical flow of image descriptors across successive frames to estimation camera motion. The fruit tracking algorithm uses the estimated camera motion between frames to predict the locations of fruits detected in previous frames. These detections are compared with fruits detected in current frame to ensure previously detected fruits are not recounted.

Using our fruit detection and tracking approach, we carried out automated counting of a block of orange trees consisting of ~ 890 trees distributed along 26 rows. Since many fruits are occluded from the camera during scans, we needed a model to predict the total count of fruits on trees based on fruits observable by the camera. To develop this model, we carried out manual fruit count of 26 trees distributed uniformly in the block. We learned a linear model that maps sensor observed fruit count to the actual count. This model was used to estimate the true yield for the whole block. Figure 15 shows a 3-D yield map, with color representing the fruit count per 0.12m along the row of orange trees. This figure is useful in comparing the overall Using our approach, we computed the total number of oranges for the 180m x 180m block to be 479,395 or on average 539 fruits per tree.

A limitation of our automated fruit counting algorithm is its dependence on suitable illumination. For example, fruit



(a) Fruits are detected using an SVM pixel classifier along with additional logic to count fruits in clusters. Blue boxes show blob boundaries, and red circles show detected fruits within the blobs.



(b) Fruit detection and tracking is carried out simultaneously to generate a running count of fruits along a row. The red boxes show detection of new fruits, the cyan boxes show new fruits that were counted in the frame. The cyan lines show the paths of a detected fruit from previous images into the current image. The purple boxes show predicted positions of the fruits detected in previous frames.

Fig. 14: Fruit detection and tracking.




Fig. 15: A 3-D map of orange trees in a 180m x 180m block in Orange Cove, CA, with the color of the point cloud depicting fruit density. Warmer colors represent higher yield. The total number of oranges for the block was estimated to be 479,395. On average, there were 539 fruits per tree.

detection quality is poor when the sun faces the camera. Additionally, color based fruit classifiers are susceptible to false positives due to canopy features that sometimes match fruit features. For example, detection and counting of green oranges is challenging due to the color similarities with canopy. To address this issue, we are exploring the use of thermal imagery along with visible camera data to exploit the temperature differential between fruits and canopy in order to improve detection accuracy. Figure 16) illustrates canopy and fruit temperature differential for apples at the orchard in Biglerville, Pennsylvania.

IV. CONCLUSIONS AND FUTURE DIRECTIONS

This paper presented a methodology for precision agriculture consisting of a light-weight multi-spectral 3-D imaging system capable of deployment onboard low-flying micro aerial vehicles, as well as a shoulder mounted harness. The hardware design of the sensor suite, and the underlying software stack was described. We developed data analysis and visualization methodologies to help growers obtain actionable information from acquired data, and interpret them for efficient farm management. Towards this end, we




Fig. 16: Thermal and visible images of apple trees captured by the sensor suite. A few apples have been highlighted in boxes, showing warmer temperature than the foliage.

demonstrated the extraction of four data products crucial to resource management – plant morphology, normalized difference vegetative index (NDVI), leaf area, and fruit count. To the best of our knowledge, our work presents the first versatile system for automated monitoring in precision agriculture, enabling sustainable production. The compact and self-contained design allows our system to be used in multiple deployment modalities, i.e., onboard UAVs and UGVs, as well as carried by human scouts. This provides

flexibility to growers to choose the appropriate mode of use. When deployed onboard multiple UAVs, our system can carry out rapid mapping of a farm, especially for applications that are time sensitive. We demonstrated the utility of our system through results from field trials at a vineyard in Galt, California, an apple orchard in Biglerville, Pennsylvania, and a citrus grove in Orange Cove, California.

Although in this work we presented UAV data from the top of canopies looking down, in our future work we will carry out UAV flights with the sensor suite pointing sideways at tree canopies, a capability currently demonstrated on the harnessed deployment mode. Exploiting the low-cost design of our sensor suite and building upon our expertise in autonomy for UAV swarms, our long-term goal is to deploy sensor suites on board multiple UAVs, facilitating rapid mapping of large farms for tasks that are time sensitive. Examples include multi-spectral imaging for monitoring crop stress, a task that demands deployment over narrow time windows due to the movement of the sun that results in changes in illumination.

ACKNOWLEDGMENTS

We want to thank Dr. Kari Peter at the Fruit Research and Extension Center (FREC) in Biglerville, Pennsylvania, Dr. Luis Sanchez at E&J Gallo Winery in Galt, California, and Eric Hoffman of Booth Ranches, California for providing access to their facilities for field experiments.

REFERENCES

- [1] M. Bryson, A. S. Reid, F. T. Ramos, and S. Sukkarieh, "Airborne Vision-based Mapping and Classification of Large Farmland Environments," *Journal of Field Robotics*, vol. 27, no. 5, pp. 632–655, 2010.
- [2] P. J. Zarco-Tejada, J. A. Berni, L. Suárez, and E. Fereres, "A new era in remote sensing of crops with unmanned robots," *SPIE Newsroom*, pp. 2–4, 2008.
- [3] U. Weiss, P. Biber, S. Laible, K. Bohlmann, and A. Zell, "Plant Species Classification Using a 3D LIDAR Sensor and Machine Learning," in *Machine Learning and Applications (ICMLA), 2010 Ninth International Conference on*, Dec 2010, pp. 339–345.
- [4] S. Singh, M. Bergerman, J. Cannons, B. Grocholsky, B. Hamner, G. Holguin, L. Hull, V. Jones, G. Kantor, H. Koselka, et al., "Comprehensive Automation for Specialty Crops: Year 1 results and lessons learned," *Intelligent Service Robotics*, vol. 3, no. 4, pp. 245–262, 2010.
- [5] T. Bakker, K. A. van, J. Bontsema, J. Miller, and G. S. van, "Systematic design of an autonomous platform for robotic weeding," *Journal of Terramechanics*, vol. 47, no. 2, pp. 63 – 73, 2010.
- [6] C. Cariou, R. Lenain, B. Thuilot, and M. Berducot, "Automatic guidance of a four-wheel-steering mobile robot for accurate field operations," *Journal of Field Robotics*, vol. 26, no. 6-7, pp. 504–518, 2009.
- [7] M. Bergerman, S. Singh, and B. Hamner, "Results with autonomous vehicles operating in specialty crops," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, May 2012, pp. 1829–1835.
- [8] J. Arn, A. Escol, J. Valls, J. Llorens, R. Sanz, J. Masip, J. Palacn, and J. Rosell-Polo, "Leaf area index estimation in vineyards using a ground-based LiDAR scanner," *Precision Agriculture*, vol. 14, no. 3, pp. 290–306, 2013.
- [9] P. Tokekar, J. Vander Hook, D. Mulla, and V. Isler, "Sensor planning for a symbiotic uav and ugv system for precision agriculture," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 5321–5326.
- [10] O. W. Liew, P. C. J. Chong, B. Li, and A. K. Asundi, "Signature optical cues: emerging technologies for monitoring plant health," *Sensors*, vol. 8, no. 5, pp. 3205–3239, 2008.
- [11] J. Berni, P. Zarco-Tejada, L. Suarez, and E. Fereres, "Thermal and Narrowband Multispectral Remote Sensing for Vegetation Monitoring From an Unmanned Aerial Vehicle," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 47, no. 3, pp. 722–738, March 2009.
- [12] A. Jaakkola, J. Hyyp, A. Kukko, X. Yu, H. Kaartinen, M. Lehtomki, and Y. Lin, "A low-cost multi-sensoral mobile mapping system and its feasibility for tree measurements," *{ISPRS} Journal of Photogrammetry and Remote Sensing*, vol. 65, no. 6, pp. 514 – 522, 2010.
- [13] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, vol. 3, no. 3.2, 2009, p. 5.
- [14] J. Sola, "Quaternion Kinematics for the error-state KF," February 2014. [Online]. Available: <http://www.iri.upc.edu/people/jsola/JoanSola/objectes/notes/kinematics.pdf>
- [15] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013.
- [16] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, Nov 2013, pp. 1280–1286.
- [17] P. Furgale, T. Barfoot, and G. Sibley, "Continuous-time batch estimation using temporal basis functions," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, May 2012, pp. 2088–2095.
- [18] R. Gade and T. B. Moeslund, "Thermal cameras and applications: A survey," *Mach. Vision Appl.*, vol. 25, no. 1, pp. 245–262, Jan. 2014. [Online]. Available: <http://dx.doi.org/10.1007/s00138-013-0570-5>