

Multimodal Dataset for Localization, Mapping and Crop Monitoring in Citrus Tree Farms

Hanzhe Teng, Yipeng Wang, Xiaao Song, and Konstantinos Karydis

University of California Riverside, Riverside, CA 92521, USA
{hteng007,ywang1040,xsong036,karydis}@ucr.edu

Abstract. In this work we introduce the CitrusFarm dataset, a comprehensive multimodal sensory dataset collected by a wheeled mobile robot operating in agricultural fields. The dataset offers stereo RGB images with depth information, as well as monochrome, near-infrared and thermal images, presenting diverse spectral responses crucial for agricultural research. Furthermore, it provides a range of navigational sensor data encompassing wheel odometry, LiDAR, inertial measurement unit (IMU), and GNSS with Real-Time Kinematic (RTK) as the centimeter-level positioning ground truth. The dataset comprises seven sequences collected in three fields of citrus trees, featuring various tree species at different growth stages, distinctive planting patterns, as well as varying daylight conditions. It spans a total operation time of 1.7 hours, covers a distance of 7.5 km, and constitutes 1.3 TB of data. We anticipate that this dataset can facilitate the development of autonomous robot systems operating in agricultural tree environments, especially for localization, mapping and crop monitoring tasks. Moreover, the rich sensing modalities offered in this dataset can also support research in a range of robotics and computer vision tasks, such as place recognition, scene understanding, object detection and segmentation, and multimodal learning. The dataset, in conjunction with related tools and resources, is made publicly available at <https://github.com/UCR-Robotics/Citrus-Farm-Dataset>.

Keywords: Datasets · Agricultural Robotics · Precision Agriculture · Crop Monitoring · Localization and Mapping · Multimodal Perception

1 Introduction

Crop monitoring is an essential and critical component in precision agriculture. Its value is that it can provide growers and agronomists with different pieces of information that are useful for determining a range of indicators, such as plant growth [10], water stress level and health condition [17], to name a few. This can lead to more informed future decisions on irrigation, disease prevention, pest control, fruit harvesting and ultimately higher yields [16]. Given the dynamic nature of changes that happen in the field, crop monitoring tasks need to be performed regularly, a process that can quickly become time- and labor-intensive and scale poorly both spatially and temporally [2]. These challenges can be

addressed to a certain extent by automating parts of the overall process with mobile robots.

Datasets are an essential tool to help develop autonomous robot systems in agricultural environments. Navigational sensor data is crucial for the development of accurate localization and mapping algorithms, which can enable a range of fully automated tasks such as soil conductivity measurement [5] and physical sample retrieval [4]. On the other hand, multispectral images are particularly valuable in agricultural research, as they can provide informative measurements for computing domain-specific indices. For example, the Normalized Difference Vegetation Index (NDVI) can be computed from red and near-infrared spectral channels and serves as an indicator of plant health [17]. If both sensing modalities are available, a detailed and accurate map with multimodal data (e.g., thermal, NDVI) can be constructed by an automated mobile robot for crop monitoring on a regular basis. However, there currently exist limited public datasets that provide *both multispectral images and navigational sensor data* to facilitate the development of such autonomous robot systems.

To this end, we present the CitrusFarm dataset, a comprehensive multi-sensor dataset in the citrus tree farms. The dataset comprises seven sequences collected in three fields of citrus trees, featuring various tree species at different growth stages, distinctive planting patterns, as well as varying daylight conditions. The dataset offers stereo RGB images with depth information, as well as monochrome, near-infrared and thermal images, presenting diverse spectral responses crucial for agricultural research. Furthermore, it provides a range of navigational sensor data encompassing wheel odometry, LiDAR, inertial measurement unit (IMU), and GNSS with Real-Time Kinematic (RTK) as the centimeter-level positioning ground truth. In comparison to related works, our CitrusFarm dataset provides a total of nine sensing modalities, thus enabling significant potential for multimodal and crossmodal research in the robotics and computer vision community.

We anticipate that this dataset can facilitate the development of autonomous robot systems operating in agricultural tree environments, primarily for localization, mapping and navigation tasks. With both multispectral and navigational data, automated multimodal mapping becomes possible for crop monitoring on a regular basis. Given the rich sensing modalities, sensor fusion can be conducted with any combination of sensing sources of interest. This can advance the research on localization in challenging unstructured agricultural environments, and enable a full range of automated agricultural operations contingent on accurate localization. Lastly, this dataset can serve as the training and testing data in machine learning tasks, and support research in a variety of robotics and computer vision applications, such as scene understanding, segmentation, object detection, place recognition and multimodal learning.

2 Related Works

Agricultural robotics datasets that provide both multispectral images and navigational sensor data are scarce. Lu and Young reviewed public agricultural image datasets for weed control, fruit detection and miscellaneous applications, without navigational sensors involved [11]. In the following, *we focus on public datasets that contain both images and navigational data*. Overall, such datasets can be categorized into two broad types.

Datasets recorded following the same route over time: Chebrolu et al. recorded the dataset on a sugar beet farm over a period of three months, aiming to advance research on plant classification, localization and mapping [6]. Bender et al. presented a multimodal dataset that contains weekly scans of cauliflower and broccoli covering a ten-week growth cycle from transplant to harvest using the Ladybird agricultural robot [3]. In addition, that dataset also provides physical characteristics of the crop, soil sensor data and environmental data from a weather station. Polvara et al. collected multi-sensor data following the same route in a vineyard environment over a time span of six months; the resulting dataset can be used for long-term localization and mapping, phenotyping and crop mapping tasks [13].

Datasets collected in a variety of scenes: Pire et al. collected six sequences of multi-sensor data on soybean fields, with the focus being on visual localization and mapping under various challenging situations [12]. Gldenring et al. focused on the problem of weed control using vision-equipped robots and presented the RumexWeeds dataset that comprises color images with annotations and navigational sensing data associated with each image frame [9]. Ali et al. recorded the dataset in a forest environment in various seasons and daylight conditions using a sensor rig mounted on a car vehicle [1].

Our work fits into the second category, and aims to diversify the data collection process in the citrus tree environments. We collected seven sequences in three agricultural fields featuring various tree species at different growth stages, distinctive planting patterns, as well as varying daylight conditions, all while following various routes in the fields. More details about our data collection process are to be discussed in Section 5.


Moreover, in comparison to related works, our CitrusFarm dataset provides a total of nine sensing modalities, thereby unlocking vast potential for multimodal research in robotics and computer vision. Among the nine modalities, the presence of LiDAR data distinguishes this work from the offering geo-referenced images alone, and multispectral images further provide with a unique perceptual perspective. A detailed overview of the sensing modalities encompassed within each dataset is presented in Table 1.

3 Mobile Robot Setup

We employ the Clearpath Jackal wheeled mobile robot to collect multi-sensor data in the agricultural fields. Jackal is a differential-drive robot that follows

Table 1. Summary of Available Sensing Modalities in Related Public Datasets and Comparison with Our Developed Dataset.

Dataset	Mono	RGB	Depth	NIR	Thermal	LiDAR	Wheel-Odom	IMU	GPS(-RTK)
Chebrolu et al. [6]		✓	✓	✓		✓	✓		✓
Bender et al. [3]	✓	✓	✓	✓	✓			✓	✓
Polvara et al. [13]		✓	✓			✓	✓	✓	✓
Pire et al. [12]		✓					✓	✓	✓
Güldenring et al. [9]		✓					✓	✓	✓
Ali et al. [1]		✓						✓	✓
Ours	✓	✓	✓	✓	✓	✓	✓	✓	✓

**Fig. 1.** (a) The Clearpath Jackal mobile robot with a variety of sensors mounted onboard, ready to collect data in the field. (b) The corresponding reference frames of five camera lenses, IMU, LiDAR and wheel odometry superimposed on the Jackal robot. Relative poses between frames are real parameters after extrinsic calibration. Red, green and blue bars denote x, y and z axes respectively.

the unicycle kinematic model, which can take linear and angular velocities as the control input. It can support in-place turns and a maximum linear speed of 2 m/s. The onboard computer installed on the robot is a Jetson Xavier AGX, with 32 GB RAM, 2 TB storage and Ubuntu 18.04 operating system. The Robot Operating System (ROS) is utilized as the middleware to connect with all sensors and record data.

On top of the robot base, we designed and manufactured several mounting parts to secure a variety of sensors onboard. Figure 3 illustrates the sensor setup on the robot. We discuss the specification and configuration of each sensing modality in detail as follows.

Monochrome Camera: We use the FLIR Blackfly BFS-U3-16S2M-CS camera to capture monochrome images. It is equipped with a Computar A4Z2812CS-MPIR lens, with an adjustable focal length and aperture. This monochrome modality can reflect the overall light intensity in the visible spectrum, and is useful for both multimodal perception and vision-based autonomous navigation (e.g., visual odometry [14]).

Stereo RGB Camera: We employ the Stereolabs Zed2i camera to capture color images in visible spectrum. This camera possesses two high-resolution RGB lenses, and depth images can be computed from stereo images in real-time by leveraging the GPU of the Xavier onboard computer. Furthermore, we also include the internal IMU data and the poses computed by the Zed camera driver in our data collection process.

Red-Green-NIR Camera: We use the Mapir Survey3 camera to capture images with three channels: Red, Green, and Near-InfraRed (NIR). The merit of this multispectral configuration is the capability to compute domain-specific indices directly from a single image. For instance, the Normalized Difference Vegetation Index (NDVI) can be derived from the Red and NIR channels, via the equation $NDVI = \frac{NIR - Red}{NIR + Red}$. Such indices can enable growers and agronomists to keep track of the health status of plants (e.g., [17]).

Thermal Camera: We consider the FLIR ADK thermal camera to capture images in the Long-Wave InfraRed (LWIR) range (around 8-14 micrometers). This thermal modality can effectively detect humans and animals, as well as facilitate the understanding of plant health condition (e.g., water stress level, disease prevention, pest control). Furthermore, a thermal modality can also enhance navigation tasks in degenerated environments (e.g., [7]).

A summary of the different cameras used in our robot is provided in Table 2.

Table 2. Summary of Camera Specification and Configuration.

Camera	Modality	Shutter Rate	Resolution	HFOV	Bit/Channel
FLIR Blackfly	Monochrome	Global 10 Hz	1440 × 1080	72°	8 × 1
FLIR ADK	Thermal	Global 10 Hz	640 × 512	65°	8 × 1
Stereolabs Zed2i	Stereo RGB	Rolling 10 Hz	1280 × 720	102°	8 × 3 × 2
Stereolabs Zed2i	Depth	Rolling 10 Hz	1280 × 720	102°	32 × 1
Mapir Survey3	Red-Green-NIR	Rolling 10 Hz	1280 × 720	85°	8 × 3

Note: Although most cameras can support up to 30 Hz or 60 Hz frame rate, we operate them at 10 Hz to match the LiDAR’s operating frequency.

3D LiDAR: We utilize a Velodyne VLP-16 LiDAR for 3D mapping. It comprises 16 scan lines (constituting 30° vertical FOV) and each line contains 1800 points (constituting 360° horizontal FOV), providing a total of 28800 points per point cloud at a rate of 10 Hz. The maximum measurement distance extends to 100 m, with an accuracy of ±3 cm.

IMU: A MicroStrain 3DM-GX5-AHRS IMU is installed underneath the LiDAR to collect linear acceleration and angular velocity data for inertial navigation. It comes equipped with built-in complementary filter and Extended Kalman Filter (EKF) algorithms that operate ahead of output, at a maximum frequency of 1000 Hz. In our configuration, we adjust its rate to 200 Hz. The internal reference frames of the accelerometer and gyroscope are treated as identical.

Wheel Odometry: We record the wheel odometry data provided by the Jackal mobile robot software, computed from its wheel encoders at a rate of 50 Hz. However, this measurement can drift over time and requires fusion with other sensors for accurate estimation. The reference frame of this sensing modality is located at the bottom center of the robot base.

GPS-RTK: We utilize the SwiftNav Duro GPS system to function as the positioning ground truth. By leveraging Real Time Kinematic (RTK) technology, it is possible to acquire centimeter-level positioning accuracy of the robot at a rate of 10 Hz. The raw measurements, denoted as (latitude, longitude, altitude), are transformed to a path aligned with the origin as per the WGS84 standard, utilizing a Python script.

4 Sensor Calibration

The calibration of all sensors is conducted in three steps: 1) calibration of both intrinsic and extrinsic parameters of all cameras, 2) calibration of the extrinsic parameter between the IMU and one selected camera, and 3) calibration of the extrinsic parameter between the LiDAR and one selected camera. We elaborate on each step in the following subsections, and provide a summary discussion at the end of this section.




Fig. 2. Sample images of the calibration target with checkerboard pattern in (a) monochrome modality, (b) thermal modality, and (c) thermal modality with grayscale values inverted. Images are captured under direct sunlight.

4.1 Multi-Camera Calibration

We use the Kalibr [8] toolbox to calibrate four cameras at once. A 24-inch square calibration target with a checkerboard pattern (Figure 2(a)) is placed at multiple

locations under direct sunlight, to cover most of the field of view of all cameras. Pinhole camera model and radtan (radial-tangential) distortion model are applied in the optimization. Results comprise the intrinsics (including distortion coefficients) of each camera lens, and a chain of transformation matrices from the first to the last camera lens.

One notable challenge we addressed in this process is the recognition of the checkerboard pattern in the thermal modality. In a typical indoor environment, the checkerboard pattern cannot be captured by the thermal camera, since the temperature of the calibration target at each point is almost the same. Considering that black areas absorb more heat than white areas, we placed the calibration target outdoor under direct sunlight instead. The resulting pattern being captured by the thermal camera is illustrated in Figure 2(b). We then invert its grayscale values to make it consistent with other checkerboard patterns for use in the Kalibr toolbox (see Figure 2(c)).

During calibration, the checkerboard pattern in thermal images may become blurry as the heat transfers from black cells to white cells. To alleviate this effect, a foam board was placed in between the poster paper (with the checkerboard pattern printed on it) and the medium density fiberboard (MDF) in the back to reduce heat transmission. Furthermore, the calibration can be conducted in the afternoon when sunlight is not very intense, and a fan can be used to cool down the calibration target when needed.

4.2 IMU-Camera Calibration

We use again the Kalibr [8] toolbox to calibrate the extrinsic parameter for the IMU and the FLIR Blackfly monochrome camera. A 24-inch square calibration target with an aprilgrid pattern is placed on the ground statically. With the pattern being captured by the camera all the time, we first collected a recording of about ten seconds when the robot is static, and then moved the robot (with both sensors rigidly mounted) towards x, y, z, roll, pitch and yaw directions, respectively, each lasting a few seconds. The final transformation matrix is obtained by optimizing the joint observations from the IMU and the camera.

In this process, the intrinsics of both sensors are required. We calibrate the noise density and random walk of the IMU by the Allan variance method beforehand.¹ The intrinsic parameters of the camera are taken from the results of multi-camera calibration.

4.3 LiDAR-Camera Calibration

We use the ACFR LiDAR-Camera calibration toolbox [15] to estimate the transformation between the LiDAR and the FLIR Blackfly monochrome camera. The 24-inch square calibration target with a checkerboard pattern is reused herein, and is placed at multiple locations within the field of view of both sensors. One additional requirement is that the target has to be placed in a tilted position

¹ https://github.com/ori-drs/allan_variance_ros

(e.g., 45-degree) in mid-air, such that the LiDAR point cloud of this target can be segmented out and square edges can be detected. In this way, the calibration program can align the center of the calibration target in both modalities.

Table 3. Summary of Calibrated Extrinsic Parameters.

Parent Frame	Child Frame	x [m]	y [m]	z [m]	q_x	q_y	q_z	q_w
base_link	velodyne_lidar	0.0400	0.0000	0.3787	0.0000	0.0000	0.0000	1.0000
velodyne_lidar	gps_rtk	-0.2200	0.0000	0.1530	0.0000	0.0000	0.0000	1.0000
velodyne_lidar	flir_blackfly	0.2178	0.0049	-0.0645	0.5076	-0.4989	0.4960	-0.4974
flir_blackfly	imu_link	-0.0061	0.0157	-0.1895	0.4987	0.5050	-0.4987	0.4977
flir_blackfly	zed2i_rgb_left	-0.0663	0.0956	-0.0161	0.0020	-0.0081	0.0031	1.0000
zed2i_rgb_left	zed2i_rgb_right	0.1198	-0.0003	-0.0046	0.0013	0.0013	0.0000	1.0000
zed2i_rgb_right	flir_adk	0.0251	-0.0948	-0.0203	0.0026	-0.0032	0.0059	1.0000
flir_adk	mapir	-0.1608	-0.0046	-0.0138	0.0028	0.0186	-0.0094	0.9998

4.4 Summary

After completing the calibration in these three steps, we can obtain the extrinsic parameters for seven modalities. The remaining two modalities are wheel odometry and GPS-RTK. Their extrinsic parameters are represented with respect to the LiDAR sensor, and can be measured and computed from the computer-aided design (CAD) models of the robot base, LiDAR and our mounting plates.

A summary of calibrated extrinsic parameters of all sensors is available in Table 3. The seven parameters detailed as $(x, y, z, q_x, q_y, q_z, q_w)$ indicate the translation measured in meters and the orientation presented in quaternions. Collectively, these variables define the transformation of a child reference frame with respect to a parent reference frame. The corresponding visualization of all sensor frames using these extrinsic parameters is shown in Figure 3(b).

5 Data Collection in Citrus Tree Farms

We collected the CitrusFarm dataset in the Agricultural Experimental Station at the University of California, Riverside in the summer of 2023. Three fields of citrus trees were selected in the experiments, owing to their variety in tree species, growth stages and planting patterns. Composite sample images of the three fields are shown in Figure 3.

During data collection in each field, the robot was remotely piloted and followed by a human operator to travel through various paths. The GPS-RTK base station, affixed to a tripod, was positioned in the center of the field to maintain proximity to the robot. By leveraging high-gain antennas for both the base station and the rover receiver on the robot, the reliable communication range can be increased to 150 m, with trees interposed. Notably, no instances of signal loss were observed throughout the experiments.



Fig. 3. The composite satellite image and sample photos of the three fields of citrus trees, covering various planting environments.

Table 4. Key Characteristics of Our Developed Dataset.

Field	Planting Pattern	Seq.	Distance (m)	Duration (min:sec)	Time of Day	Motion	Loop or Retrace
One	sparse, uniform	01	1533.96	21:26	3 pm	lawn-mower	No
		02	524.60	07:24	5 pm	every-fifth-row	No
		03	2035.21	28:27	7 pm	every-other-row	Yes
Two	semi-dense, in-row	04	896.15	12:49	4 pm	square-spiral	Yes
		05	865.33	12:39	6 pm	lawn-mower	Yes
Three	dense, in-row	06	356.57	05:02	4 pm	every-other-row	No
		07	1212.74	16:53	6 pm	lawn-mower	Yes

In total, we collected seven sequences in the three fields. These sequences span a total operation time of 1.7 hours, cover a total distance of 7.5 km, and constitute a total of 1.3 TB of data. Key characteristics of the dataset are summarized in Table 4. The dataset offers a total of nine sensing modalities. Figure 4 depicts sample images from five camera modalities, and one instance of a domain-specific index computed from multispectral images (i.e. NDVI). Figure 5 provides sample trajectories of selected sequences.




Fig. 4. Sample images collected in the three fields in various modalities, where depth images are computed from stereo RGB images, and NDVI images are computed from the Red and NIR channels of the Red-Green-NIR images.

The primary data format we used in data collection is ROS bags. To simplify data storage and transfer, we split recorded data into blocks of 4GB, and categorized them based on their respective modalities. This categorization allows users to selectively download the subsets that are of their interest. To utilize the dataset, users need to only place the downloaded subsets in a common folder and initiate data playback; ROS will then automatically arrange the data across all bags and sequence the playback according to timestamps. Additional details on how to utilize the dataset, along with some highlighted use cases, are provided in the open-source repository for this work.

Further, to accommodate diverse use cases, we also provide Python scripts to extract data from ROS bags into files, organized in different levels of folders. These comprise images of five modalities, point clouds obtained from 3D LiDAR, CSV files for IMU readings, wheel odometry readings, and ground-truth paths




Fig. 5. Sample point cloud maps reconstructed from 3D LiDAR data, superimposed with GPS-RTK trajectories, for (a) sequence 03, (b) sequence 05, and (c) sequence 06 (sequence details are listed in Table 4).

from the GPS-RTK data. Moreover, raw data collected during the calibration process and the computed calibration parameters are included in the dataset.

6 Conclusion

In this work we presented the CitrusFarm dataset, a comprehensive multi-sensor data collection in the citrus tree farms. The dataset comprises seven sequences collected in three fields of citrus trees, featuring various tree species at different growth stages, distinctive planting patterns, as well as varying daylight conditions. Furthermore, it provides a total of nine sensing modalities: monochrome, stereo RGB, depth, near-infrared and thermal images, as well as wheel odometry, LiDAR, IMU and GPS-RTK data.

We anticipate that this dataset can facilitate the development of autonomous robot systems operating in agricultural tree environments, primarily for localization, mapping and navigation tasks. The rich sensing modalities provided in this dataset can promote a wide range of research in robotics and computer vision, such as sensor fusion and multimodal machine learning.

Future works based on this dataset can include the development of a multimodal mapping framework for crop monitoring, sensor fusion for localization in unstructured agricultural tree environments, and the integration of localization and mapping algorithms into fully automated agricultural operations.

Acknowledgments

We gratefully acknowledge the support of NSF # CMMI-2046270, USDA-NIFA # 2021-67022-33453, ONR # N00014-18-1-2252 and Univ. of California UC-MRPI # M21PR3417. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

1. Ali, I., Durmush, A., Suominen, O., Yli-Hietanen, J., Peltonen, S., Collin, J., Gotchev, A.: Finnforest dataset: A forest landscape for visual slam. *Robotics and Autonomous Systems* **132**, 103610 (2020)

2. Bechar, A., Vigneault, C.: Agricultural robots for field operations: Concepts and components. *Biosystems Engineering* **149**, 94–111 (2016)
3. Bender, A., Whelan, B., Sukkarieh, S.: A high-resolution, multimodal data set for agricultural robotics: A ladybird’s-eye view of brassica. *Journal of Field Robotics* **37**(1), 73–96 (2020)
4. Campbell, M., Dechemi, A., Karydis, K.: An Integrated Actuation-Perception Framework for Robotic Leaf Retrieval: Detection, Localization, and Cutting. In: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. pp. 9210–9216 (2022)
5. Campbell, M., Ye, K., Scudiero, E., Karydis, K.: A Portable Agricultural Robot for Continuous Apparent Soil Electrical Conductivity Measurements to Improve Irrigation Practices. In: *IEEE Int. Conf. on Automation Science and Engineering (CASE)*. pp. 2228–2234 (2021)
6. Chebroly, N., Lottes, P., Schaefer, A., Winterhalter, W., Burgard, W., Stachniss, C.: Agricultural robot dataset for plant classification, localization and mapping on sugar beet fields. *The International Journal of Robotics Research* **36**(10), 1045–1052 (2017)
7. Chen, L., Sun, L., Yang, T., Fan, L., Huang, K., Xuanyuan, Z.: RGB-T SLAM: A flexible SLAM framework by combining appearance and thermal information. In: *IEEE Int. Conf. on Robotics and Automation (ICRA)*. pp. 5682–5687 (2017)
8. Furgale, P., Rehder, J., Siegwart, R.: Unified temporal and spatial calibration for multi-sensor systems. In: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. pp. 1280–1286 (2013)
9. G ldenring, R., Van Evert, F.K., Nalpantidis, L.: Rumexweeds: A grassland dataset for agricultural robotics. *Journal of Field Robotics* (2023)
10. Kim, W.S., Lee, D.H., Kim, Y.J., Kim, T., Lee, W.S., Choi, C.H.: Stereo-vision-based crop height estimation for agricultural robots. *Computers and Electronics in Agriculture* **181**, 105937 (2021)
11. Lu, Y., Young, S.: A survey of public datasets for computer vision tasks in precision agriculture. *Computers and Electronics in Agriculture* **178**, 105760 (2020)
12. Pire, T., Mujica, M., Civera, J., Kofman, E.: The rosario dataset: Multisensor data for localization and mapping in agricultural environments. *The International Journal of Robotics Research* **38**(6), 633–641 (2019)
13. Polvara, R., Molina, S., Hroob, I., Papadimitriou, A., Tsiolis, K., Giakoumis, D., Likothanassis, S., Tzovaras, D., Cielniak, G., Hanheide, M.: Bacchus Long-Term (BLT) data set: Acquisition of the agricultural multimodal BLT data set with automated robot deployment. *Journal of Field Robotics* (2023)
14. Qin, T., Li, P., Shen, S.: VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator. *IEEE Transactions on Robotics* **34**(4), 1004–1020 (2018)
15. Tsai, D., Worrall, S., Shan, M., Lohr, A., Nebot, E.: Optimising the selection of samples for robust lidar camera calibration. In: *IEEE International Intelligent Transportation Systems Conference (ITSC)*. pp. 2631–2638 (2021)
16. Vougioukas, S.G.: Agricultural Robotics. *Annual Review of Control, Robotics, and Autonomous Systems* **2**(1), 365–392 (2019)
17. Zhao, T., Stark, B., Chen, Y., Ray, A.L., Doll, D.: A detailed field study of direct correlations between ground truth crop water stress and normalized difference vegetation index (NDVI) from small unmanned aerial system (sUAS). In: *IEEE Int. Conf. on Unmanned Aircraft Systems (ICUAS)*. pp. 520–525 (2015)