Horticulture
Innovation
Australia

# Final Report

## An Intelligent Farm Robot for the Vegetable Industry

Dr Salah Sukkarieh
University of Sydney

Project Number: VG12104

**VG12104**

This project has been funded by Horticulture Innovation Australia Limited using the vegetable industry levy and funds from the Australian Government.



ISBN 0 7341 3765 6

Published and distributed by:
Horticulture Innovation Australia Limited
Level 8, 1 Chifley Square
Sydney NSW 2000
Tel: (02) 8295 2300
Fax: (02) 8295 2399

# Contents

# Summary

The aim of this project was to develop and demonstrate a ground robot with supporting intelligent software that has the capability of conducting autonomous farm operations for different vegetable crop varieties.

The robotic system developed, The Ladybird Farm Robot, is a lightweight, solar-electric, and omni-directional robot, that uses intelligent perception for conducting surveillance (mapping, classification, detection) as well as providing information for future research into higher level tasks such as crop health monitoring and yield estimation. The system also has a robotic arm so that initial demonstrations of autonomous and calculated (time, date, velocity, volume, temperature, space etc.) delivery of insecticide, herbicide, fertiliser or other fluids and particulates to targets such as crop, weeds and other pests could be demonstrated. Various trials were conducted on a working farm.

An extension to the project was to develop a pre-production version of the robotic system. The project progressed through an initial design and requirements capture phase through to the design and build of a system capable of autonomously identifying and shooting at targets when stationary. This system is named RIPPA™ (Robot for Intelligent Perception and Precision Application™). The RIPPA system was initially demonstrated at the University of Sydney campus, where a demonstration and results of high speed targeted liquid shooting towards vegetation using NDVI for identification was presented. The system used for detecting and shooting targets is named VIIPA™ (Variable Injection Intelligent Precision Applicator™), and is mounted onto the RIPPA robot for all experiments described throughout this report. For each experiment, VIIPA sprays water, however as future experimentation shifts, the water can be replaced with other fluid such as herbicide, insecticide, fertiliser or air, that may also be mixed with other particulates. Tests were also conducted on a working farm.

Overall the project was highly successful demonstrating how modular low-cost technologies could come together and bring automation to the row crop horticulture industry to support improvements to land and labour productivity. The robotic systems are also primed to undertake the next phase of research, to demonstrate the capability of real time decision making to improve crop productivity.

# Keywords

Autonomous farming; robotic farming; robotic weeding; autonomous weeding; intelligent farm robot; robotic thinning; robotic pest management; autonomous targeted spraying; RIPPA; VIIPA.

# Introduction

The horticulture industry is one of the fastest growing industries in agriculture; with some 30,000 businesses nationally, and a farm gate value of $9 billion. As the most labour intensive of all agricultural industries, Horticulture employs around one-third of those employed in agriculture as a whole.

The increasing costs of production associated with seasonal labour force, greater scrutiny of food safety issues and consumer expectations for environmentally responsible production processes, are driving the industry to better understand, measure and strategically respond to issues involving mechanisation and automation capability. High labour costs and lack of reliable supply of reliable labour force severely limit the export competitiveness of Australian horticulture products as well as increase the risks of increased import substitution. Mechanisation and automation are a high priority for Australian horticulture.

Perception is considered to be the key element in the development of any robotic ground vehicle that has to operate in field conditions. Multi-modal sensor combinations form the basis of these perception systems. A multi-model sensor suite is a set of sensors that operate with different modalities, usually with different physical properties. Typically operating in different regions of the electro-magnetic spectrum, (e.g. visible light, far infrared, mm-wave radar) these sensors provide rich discrimination capability and robustness through redundancy, while avoiding common mode failure. Perception solutions include the ability for a robot to provide centimetre-mapping accuracy, to be able to classify features within this map (i.e. to determine and class what is being seen), to detect features of importance (such as weed, or a fruit of certain size/ripeness). With perception solved the more labour intensive tasks can be automated.

The main system components and elements of the project were:

- Robotic Platform – the design, development and testing of a robotic platform capable of traversing through crop rows with appropriate sensing comprising of spectral, vision and laser sensor. The platform was to be both remote controlled and have autonomous traversing capability. All data is to be logged on board with computing, communication and battery requirements. The system had to have solar recharge capability on board.

- Algorithms – the main research and development was to focus on map generation and visualisation; change detection, machine learning for crop, and weed identification; and platform control and path planning. The algorithm activity would be a constant evolution throughout the life of the project.

- Demonstrations – the field trials were to be conducted on a working farm to demonstrate the capabilities of the system traversing over actual farm terrain, and dealing with various aspects such as weather and changing lighting conditions.

- Where possible to focus on commercial pathways for the various technologies.

# Methodology

## *The Ladybird Farm Robot*

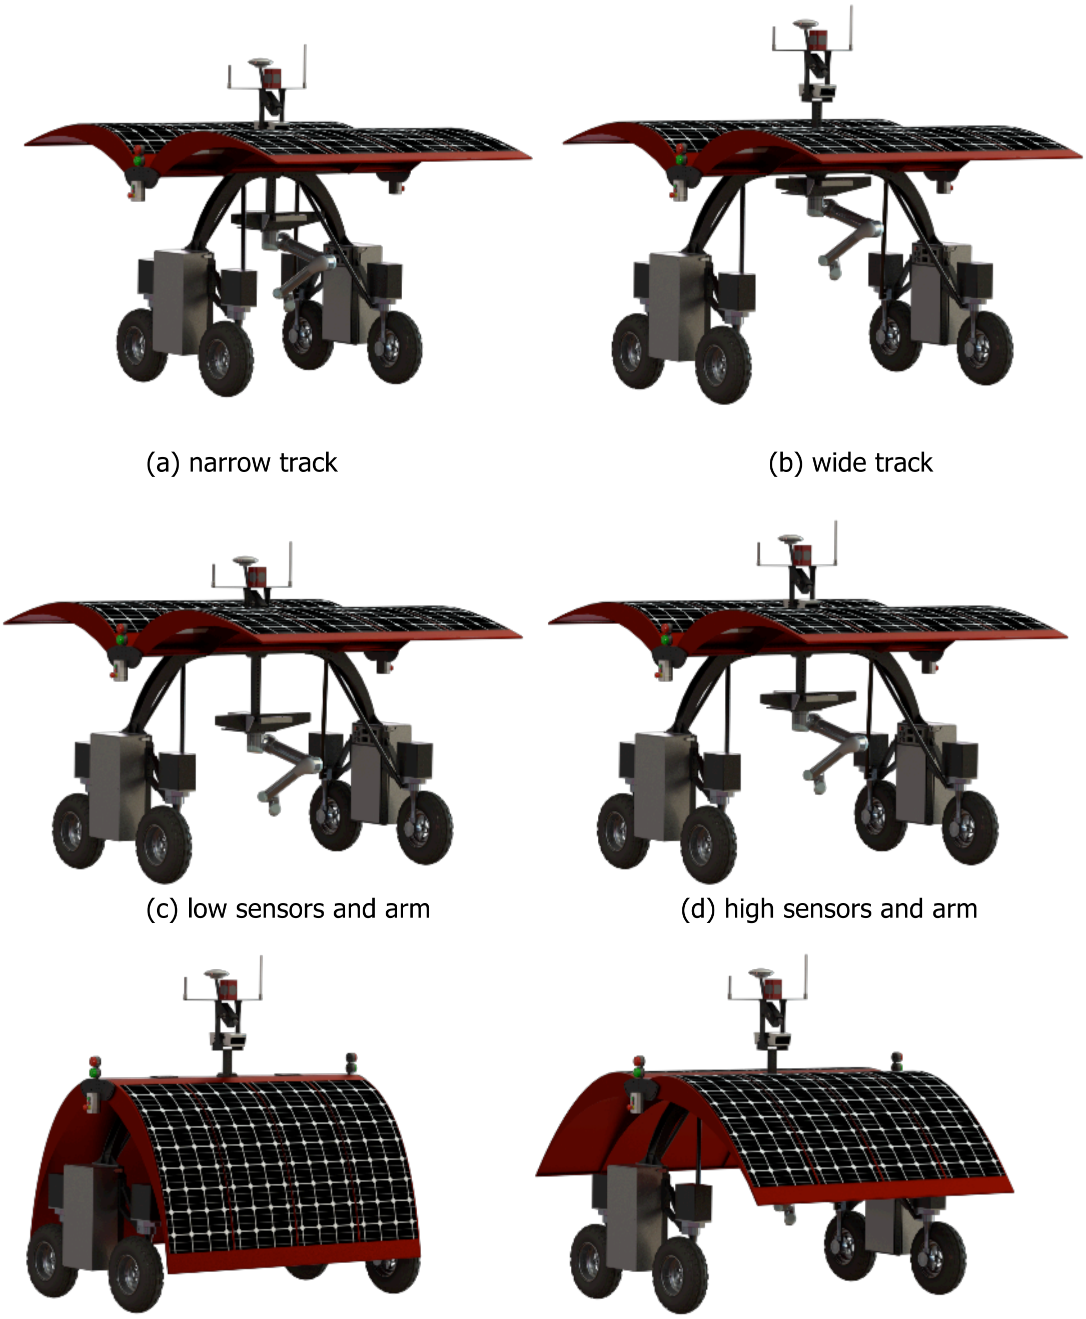

(a) narrow track (b) wide track

(c) low sensors and arm (d) high sensors and arm

(e) low cover (f) high cov

**Figure 1: Static configurations of the Ladybird.**

## Static Chassis Configuration

The geometry of the Ladybird chassis was designed to be adjustable, to match a wide variety of typical vegetable farm configurations. Three main adjustments can be made: the width of the wheelbase, the height of the sensor-manipulator payload, and the angle of the covers, as illustrated in Fig. 1. The track width can be continuously adjusted between 1.5 and 2.3m, to suit standard tractor widths found on typical farms (or any non-standard width in this range), and the height of the sensor/manipulator payload is continually adjustable over a range of 0.6m to suit different crop heights.

## Dynamic Wheel Configuration

The Ladybird utilises an omni-directional wheel configuration. There are four independent wheel modules and subsystems, one at each corner of the Ladybird. Each module is capable of driving the wheel (with one motor), and also rotating the wheel orientation continuously (with a second motor) to any angle, including infinite rotation in either direction. Software is responsible for rotating and driving all four wheels to achieve the desired high-level motion of the Ladybird.

In all situations, the four wheels must be oriented such that the perpendicular axes align to a common focal point, which forms the centre of rotation for the whole platform. The software interprets high-level platform commands in the form of {velocity, heading, turn- rate}, calculates the correct wheel angles and rates of rotation for the wheels, and executes these commands. The result is that the vehicle is highly manoeuvrable. It is capable of driving forwards or backwards in a straight line, driving at any arbitrary angle (e.g. sideways or "crabbing"), and turning on the spot. Furthermore, it is capable of any combination of these primitives, such as driving forwards, sideways and rotating all at the same time. The primitives illustrate the flexibility of motion of the platform, and additionally there is a continuum of motion between all of the configurations shown and combinations of these primitives may be superimposed.

## Control Interfaces

The platform is controlled via a software interface of the form {velocity, heading, turn- rate}. The velocity specifies the speed at which the vehicle will move. The heading specifies the direction the vehicle will face with respect to the direction of motion. For example, a heading of 0 degrees means the vehicle moves in the direction it faces. A heading of 90 degrees means the vehicle will move sideways. Finally, the turn-rate specifies the rotational speed of the whole platform.

*Figure 2: The Ladybird scanning a field of beetroot at Mulyan Farms in Cowra, NSW, during its first field trial in June 2014.*


*Figure 3: The Hetronic remote control interface.*

## *Performance in the Field*

The Ladybird performed well in the field and the commissioning was successful. This section presents the main issues that were observed in the field by the operators.

It was easy for the operators to manually drive the vehicle around Mulyan Farms (Fig. 2), including manually driving between fields, along rows of beetroot, spinach and onions and switching between rows at the headland. The operators observed that it was easier than expected to manually drive along rows, because the trench lines between rows allow gravity to pull all of the wheels into line without any specific user command.

The standard method for manual operator control is to use the Hetronic remote control unit shown in Fig. 3. All vehicle motion is controlled by the single joystick to the right of the controller, which has three axes. The combination of up-down / left-right controls the velocity and heading of the vehicle, and twisting the joystick controls the turn-rate. The red button in the top-centre is a dedicated emergency stop button. When pressed, all moving components in the system are powered down. This is implemented in hardware and is not dependent on the state of the software. It is also possible to use a laptop computer to control the vehicle, by pressing keys that are mapped to velocity, heading and turn-rate.

It was relatively easy to pack and unpack the Ladybird from the trailer.

For most activities, during sunny days, it was observed that the Ladybird had more charge by the end of the day than it had in the morning/ The rows were oriented approximately towards the sun for most of our experimentation; when idling, we oriented the solar panels to point to the sun and although all the Ladybird subsystems (computers, sensors, etc.) were running all day, the vehicle was driven at a relatively low duty cycle. It is reasonable to expect that continuous driving to scan entire fields will drain the batteries, however, the amount of charge obtained from the sun was significant and could be used to charge the system for intermittent scans.

### *Field Trials and Data Collection*

The primary objectives of the field trials were to:

1. Commission the Ladybird robot on an exemplar commercial vegetable farm
2. Test manual and autonomous control on the farm, at three sites
3. Obtain initial camera and lidar (perception)  datasets from three sites

## Location

The field trials were conducted at Mulyan Farms near Cowra, New South Wales (Fig. 4). A large variety of crops are grown there and in discussion with the farm manager, varieties were selected based on logistics, the time of season for each crop, and to have some variability in the nature of the crops.

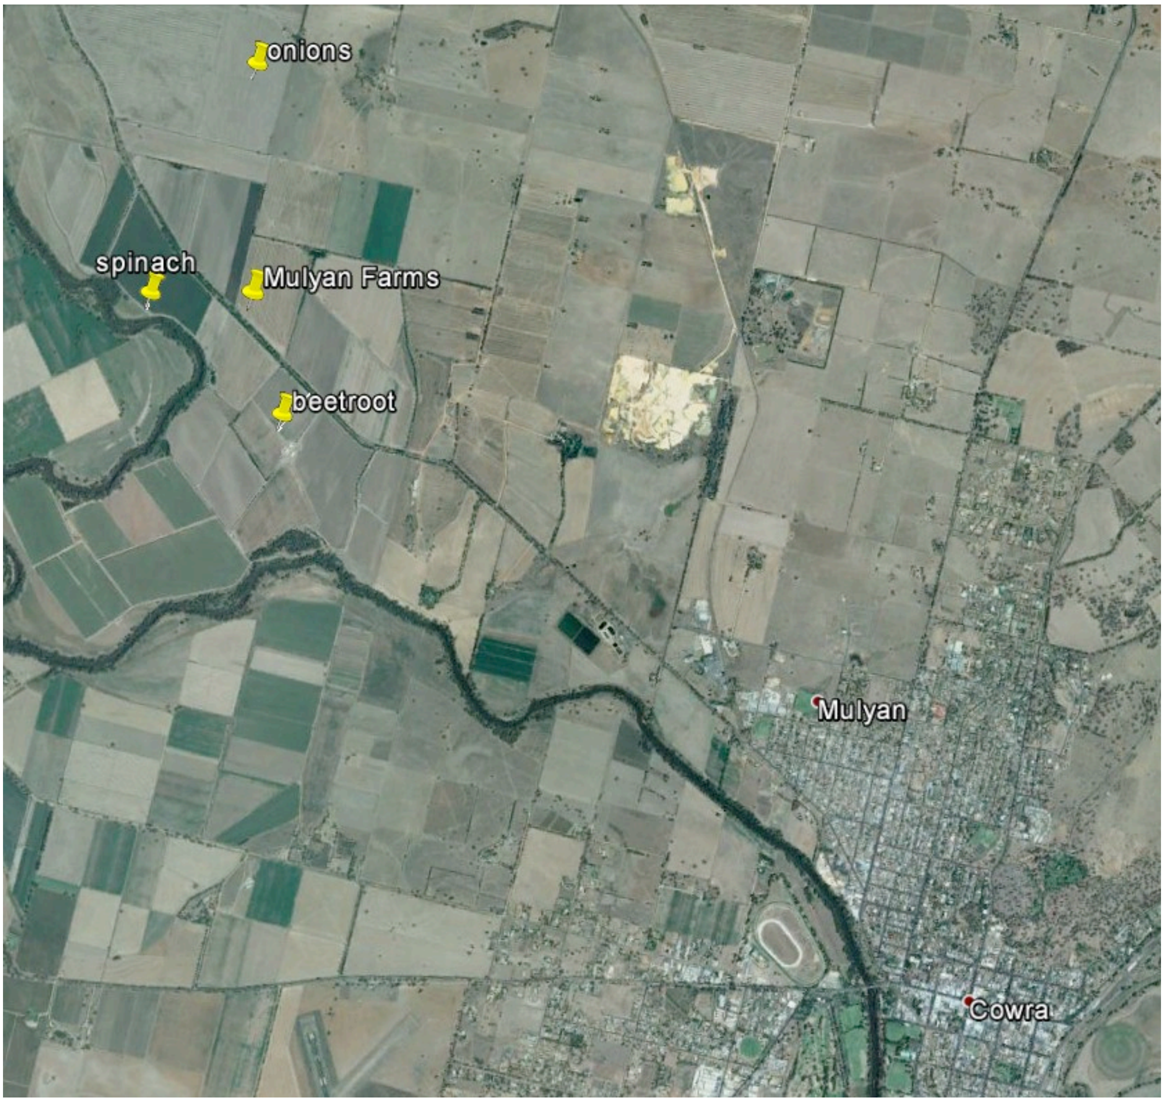


***Figure 4: Mulyan Farms, near Cowra, New South Wales. The site of the Ladybird field trials.***

Specific trajectories performed by the Ladybird can be generated from the user interface, which automatically transforms data logged from the Ladybird to a Google Earth and Google Maps compatible format (Fig. 5). Fig. 6 shows the Ladybird at the field.


*Figure 5: Areas scanned by the Ladybird. Autonomous trajectories shown in white.*

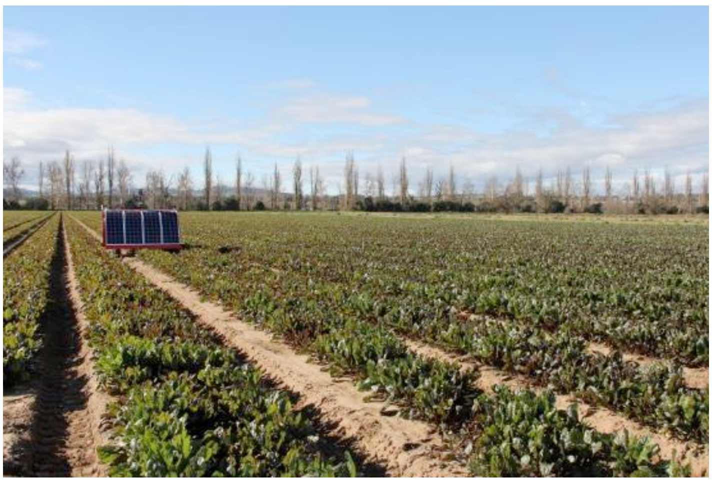

*(a) beetroot*

*(b) spinach*

*(c) onions*

*Figure 6: The Ladybird in operation at Mulyan Farms.*

## Data

Data from the Ladybird sensors were collected during all of the trajectories. The sensors were configured as shown in Fig. 7 and an example of the raw sensor data is shown for beetroot in Fig. 8. The panospheric camera provides 5 colour images looking around the vehicle and one pointing directly up. The stereo camera provides three colour images offset laterally by 10cm (one image shown only), looking directly down from underneath the Ladybird.

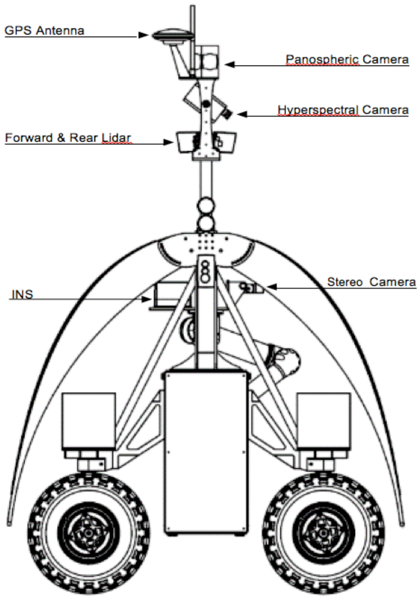


(a) sensors

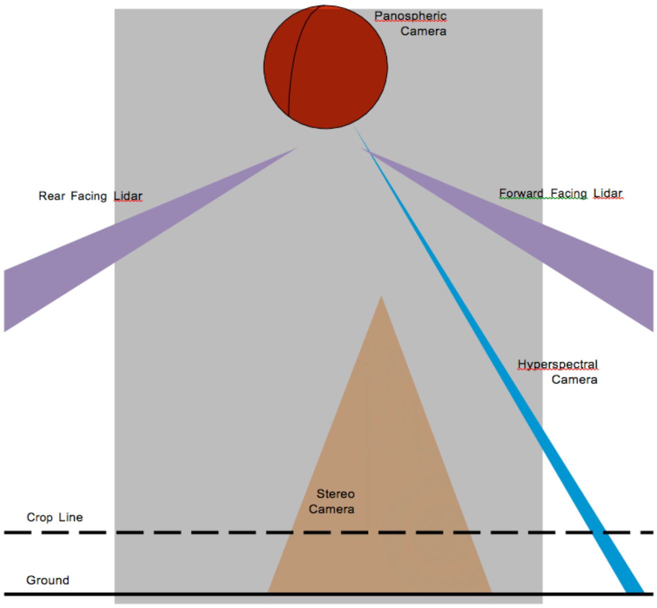


(b) fields of view

*Figure 7: Ladybird sensor configuration.*

The hyper-spectral camera provides a line scan that cuts across the vegetable row immediately in front of the Ladybird. The image shown in Fig. 8(c) is interpreted as follows: the horizontal (longer) axis is spatial, representing the swathe across the vegetable patch where the field of view of the camera intersects the ground. Pixels further to the left of the image are from one side of the vegetable row, centre pixels are from the middle of the row, and pixels to the right of the image are from the other side of the row. The vertical (shorter) axis represents the spectral dimension. From the top of the image to the bottom, pixels correspond to the continuum from blue then green, to red then infra-red respectively, divided into 240 discrete bands. The top-most pixel represents the intensity of light at 400nm and the bottom most pixels represent the intensity at 900nm. Each vertical column of pixels is the spectral response or signature from one point on the farm. As an example, the medium- bright pixels in the middle band and the brighter pixels in the band across the bottom centre are the green and near-infra-red components of the spectrum, which is the characteristic spectral signature due to chlorophyll.

Fig. 8(d) shows one scan from the forward facing lidar. The lidar range scan intersects the vegetable row with four separate beams, each at a slightly different angle, shown in a different colour in the figure. The horizontal axis is the left-to-right position across the scene in front of the vehicle, and the vertical axis is the range. As a result, trenches appear as peaks (longer range from the vehicle) and the beds of beetroot appear as shorter range segments. The data appears to be noisy, which is expected given the pseudo-random geometry of the beetroot leaves, and is also in part due to intrinsic noise from the sensor.

(a) stereo camera

(b) panospheric camera

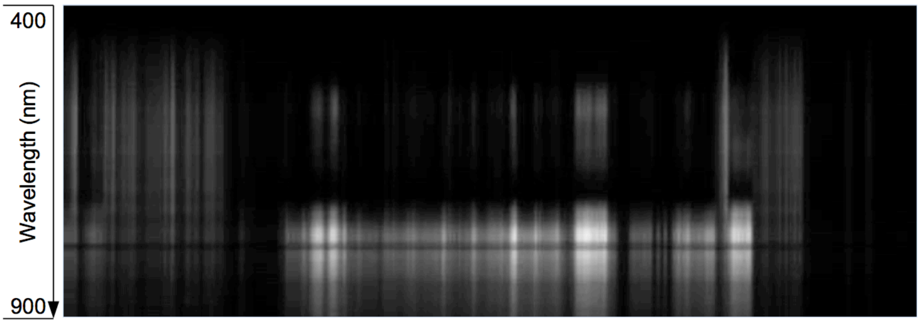


(c) hyperspectral camera

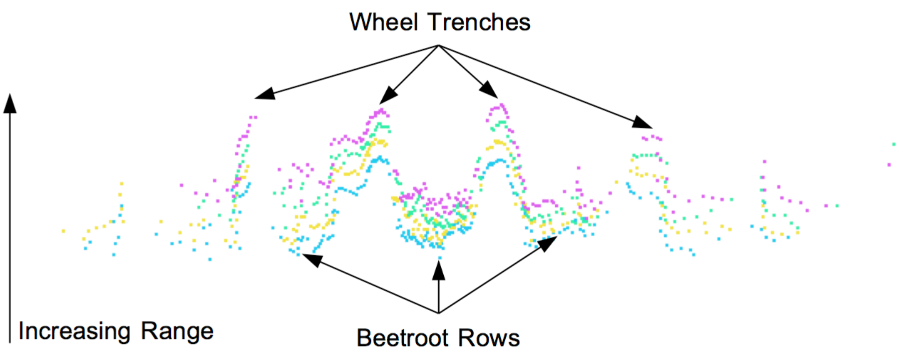


(d) front facing lidar

*Figure 8: Examples of raw sensor data, while scanning beetroot.*

## *Autonomous Operation*

Autonomous operation was achieved successfully on our first field trial at Mulyan Farms. The paddocks are laser-levelled to allow for gravity fed irrigation, and crop rows are prepared and sowed using accurate real-time kinematic, global positioning system (RTK-GPS) guided machinery. As a result, the rows are straight and spaced evenly. It is therefore possible for the Ladybird to drive autonomously by following pre-defined GPS way-points or way-lines, which can be geometrically defined according to the original field specification.

The Ladybird is equipped with a Novatel SPAN RTK/GPS/INS inertial navigation system, which, together with a base station for RTK corrections, provides a centimetre accurate solution. In the field, we provide the corrections via a base station in our trailer with a wifi bridge from the trailer to the robot. Other solutions are possible, such as re-using existing or commercial RTK sources that may be available on farm sites already. Good satellite visibility and a centimetre accurate position solution were observed for the full three day duration of our first field trip to Cowra.

### Defining the Paddock Geometry

For autonomous control along rows and between rows of a paddock, the geometry of the paddock must be defined in geographical coordinates. Specifically, the latitudinal and longitudinal coordinates of the start and end of each row must be obtained. In broad-acre applications where farms may stretch over the horizon, the curvature of the Earth at these scales must be considered to enable truly straight lines. However, for the relatively small sized paddocks, a flat Earth assumption can be used locally. We use Cartesian map grids to represent coordinates on the farm. To obtain the geometry of a paddock, we manually drove the Ladybird along the first row, and recorded the location from the navigation system at the start and end of the row. We did the same for the last row in the section of the paddock we intended to drive, and linearly divided all of the rows between. This is illustrated in Fig. 10. This method was effective for our first field trial, to obtain coordinates at the three paddocks where we operated. However, to cover an entire farm more efficiently, the coordinates from the precision machinery that prepared the field could be re-used, providing a navigation map to any robotic equipment to cover the whole farm.

***Figure 10: A way-line pattern covering 8 adjacent rows of a field of onions. Way-points are shown in blue and lines in red. This pattern specifies that the Ladybird should drive approximately 15m into each row before returning, moving to the next row, and repeating for 8 rows. Not shown are the additional parameters such as velocity and vehicle heading for each segment along a single way-line. This functionality is then repeated step by step for each subsequent way-line in the list, as each waypoint (line end-point) is reached.***

## Autonomous Control Algorithm

A set of waypoints in a list is used to specify the desired trajectory of the robot. The purpose of the control algorithm is to achieve this trajectory by driving and steering the four independent wheels of the Ladybird.

A standard cross-track way-line controller was developed. When the robot is far away from the line, it should drive perpendicular towards the line to reduce the cross-track error. When the vehicle is on or near the line, it should drive in the direction of the line. The behaviour is actually specified in a continuous manner between these two extremes. The Ladybird is an omni-directional platform, which has the ability to drive one way, while facing any other direction. It can even rotate the direction it is pointing and change the direction it is moving independently and at the same time. This gives a distinct advantage over skid-steered (or "Ackerman", car-type steered) platforms for way-line control; the heading controller can be decoupled from the cross-track controller. The vehicle can be independently rotated to point along the line, while the platform moves towards and then along the line, as illustrated in Fig. 11. Once the initial error is reduced to zero, then at all times the platform, and therefore the sensor payload, is stably oriented with respect to the field. Any disturbances due to terrain can be handled by a sideways "crabbing" motion, without having to rotate the platform.

An example of the controller performance is given in Fig. 12. The controller is accurate, other than the over-shoot seen at the right of the figure (his was due to the initially conservative limitations we placed on the platform deceleration, which will later adjusted). The controller was able to guide the vehicle repeatedly along 8 row sections of roughly 20m length (into the row) for three sections of the Mulyan Farms (beetroot, onions and spinach). It was also successfully tested on two full beetroot rows, of roughly 300m length.

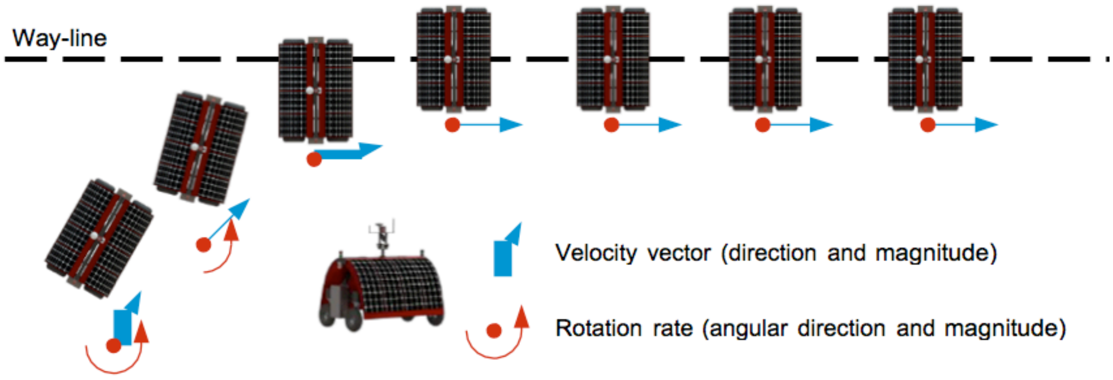


***Figure 11: A successful controller moves the vehicle in order to reduce the cross-track error to zero, while independently rotating the vehicle to reduce the heading error to zero. In this example, the cross-track error reduces from the beginning, even though the vehicle is initially pointing away from the line. The heading error also reduces immediately and in this example reaches zero before the cross-track. In the final stage, the vehicle "crabs" onto the line, reducing both errors to zero.***

***Figure 12: Actual controlled trajectory overlaid in black over the wayline. The scene is viewed from a top down perspective.***

## *Mapping, Segmentation and Classification*

In this section, we use the image data from the stereo and hyperspectral cameras to generate maps of the scanned regions of the farm. In addition, we have also performed vegetation segmentation and classification.

### Mapping

The images collected from the downward pointing stereo camera can be combined with the robot navigation data to produce a visual map of the surveyed sites. In this section the maps were generated using third party software named Agisoft (http://www.agisoft.ru/products/photoscan/professional/), which uses both images matching and the robotic navigation data to stitch the images together. The original images had slight motion blur due to the low lighting condition under the canopy of the robot, therefore it was not necessary to use the full resolution images. The images were down-sampled to speed up the process.

The map of the onion site is shown in Fig. 13. From discussions with the farm manager we learnt that the right four rows were sprayed with herbicide but not the left four, and the effect is clearly visible from the visual map. A more detailed close-up view is shown in Fig. 14. The row on the right has been sprayed, and although the onions were still in an early stage of growth, they are still visible in this map.

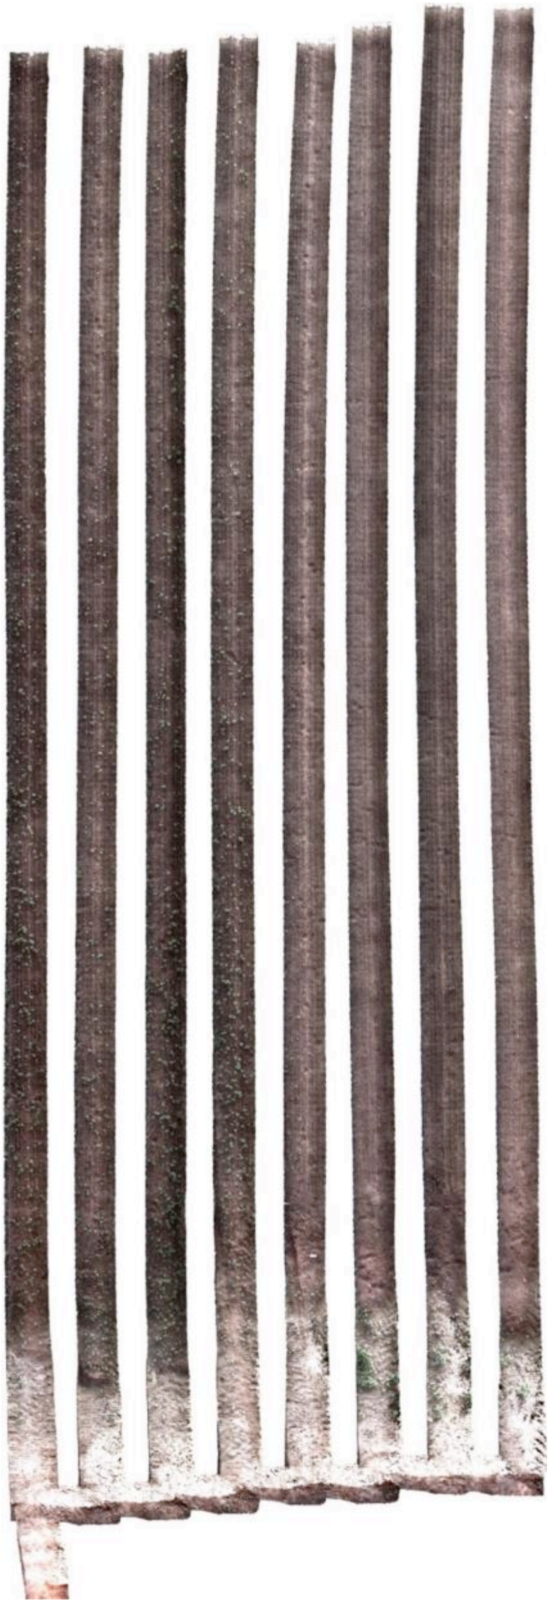

*Figure 13: The visual map (mosaic of RGB images) of the onion field.*

*Figure 14: A close-up view of the onion field.*

The visual map of spinach is shown in Fig. 15, with a close-up view in Fig. 16. It is interesting to observe that although each original image was captured in a low-light condition, the final mosaic map still contains enough information to show the different colouration of the individual spinach leaves.

The visual maps can also be viewed using Google Earth and Google Maps (Fig. 17), which allows the final map product to be easily used by the farm operators. The same interface can be used in conjunction with other data layers, such as satellite imagery, vehicle trajectories, points of interest, and ultimately (with future work) the results of high-level information processing such as yield measurements or pest and weed detection processes.

***Figure 15: The visual map (mosaic of RGB images) of the spinach site.  Note the center two rows have shorter coverage, although they were scanned by the Ladybird. The Agisoft image stitching software failure to correctly register the images in these rows, possibly due to a combination of motion blur and poor illumination, which was later addressed by adding artificial illumination to the Ladybird.***

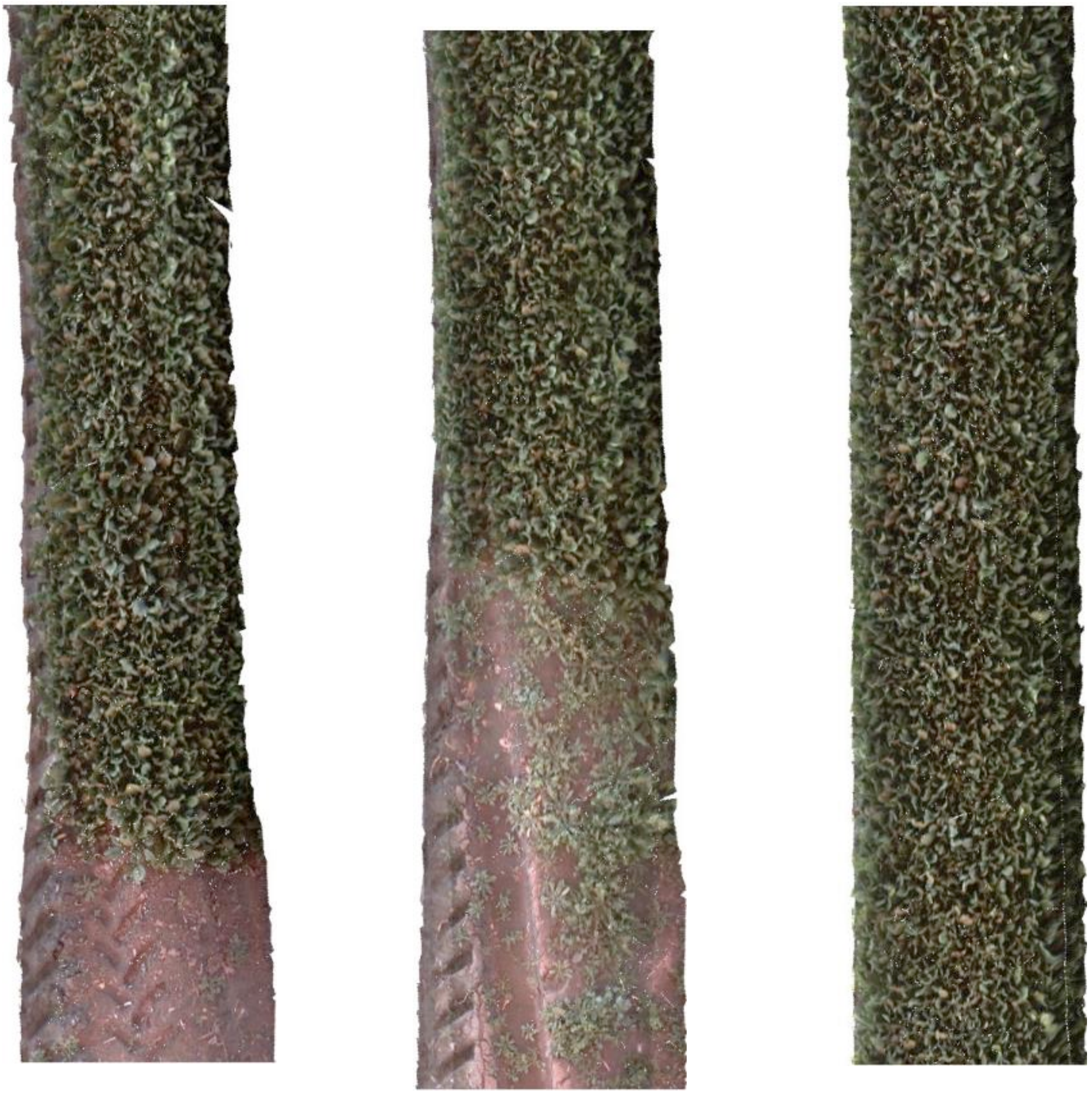

*Figure 16: A close-up view of the spinach field.*


*Figure 17: Final map product viewed on Google Earth.*

## Segmentation

In this section vegetation segmentation was performed using the data collected from the hyperspectral sensor. A common approach to vegetation segmentation is to use the Normalised Difference Vegetation Index (NDVI). Vegetation absorbs light from the visible spectrum for photosynthesis, and reflects most of the near infrared spectrum because the photon energy in this spectral region is not sufficient for photosynthesis and will only cause the plant to overheat. NDVI exploits this absorption difference between the spectral bands to identify vegetation from other non-vegetative materials. NDVI can be computed using:

$$NDVI = \frac{VIS - NIR}{VIS + NIR}$$

where VIS is the average recorded intensity in the visible spectrum, and NIR is the average in the near infrared spectrum. For satellite based remote sensing, the visible spectral range between 580nm to 680nm is used and 725nm to 900nm is used for near infrared (http://noaasis.noaa.gov/NOAASIS/ml/avhrr.htm). The hyperspectral camera on the Ladybird measures light in 240 bands between 400nm and 900nm with 240 NDVI ranges from -1 to 1. Values less than 0 correspond to water surfaces, 0.1 and below correspond to soil and rocks, and above 0.2 correspond to vegetation. In this section, values above 0.4 were used to label onion vegetation and weeds, and values above 0.6 were used to extract spinach and beetroot leaves.

An example of the hyperspectral data for (shown in RGB), the corresponding NDVI and the vegetation segmentation are shown in Figures 18-20.

This section showed that the vegetation can be successfully segmented using NDVI. However, using NDVI alone is not sufficient to distinguish between different types of vegetation (such as the crop and weeds) because they all have chlorophyll and therefore have similar NDVI values.

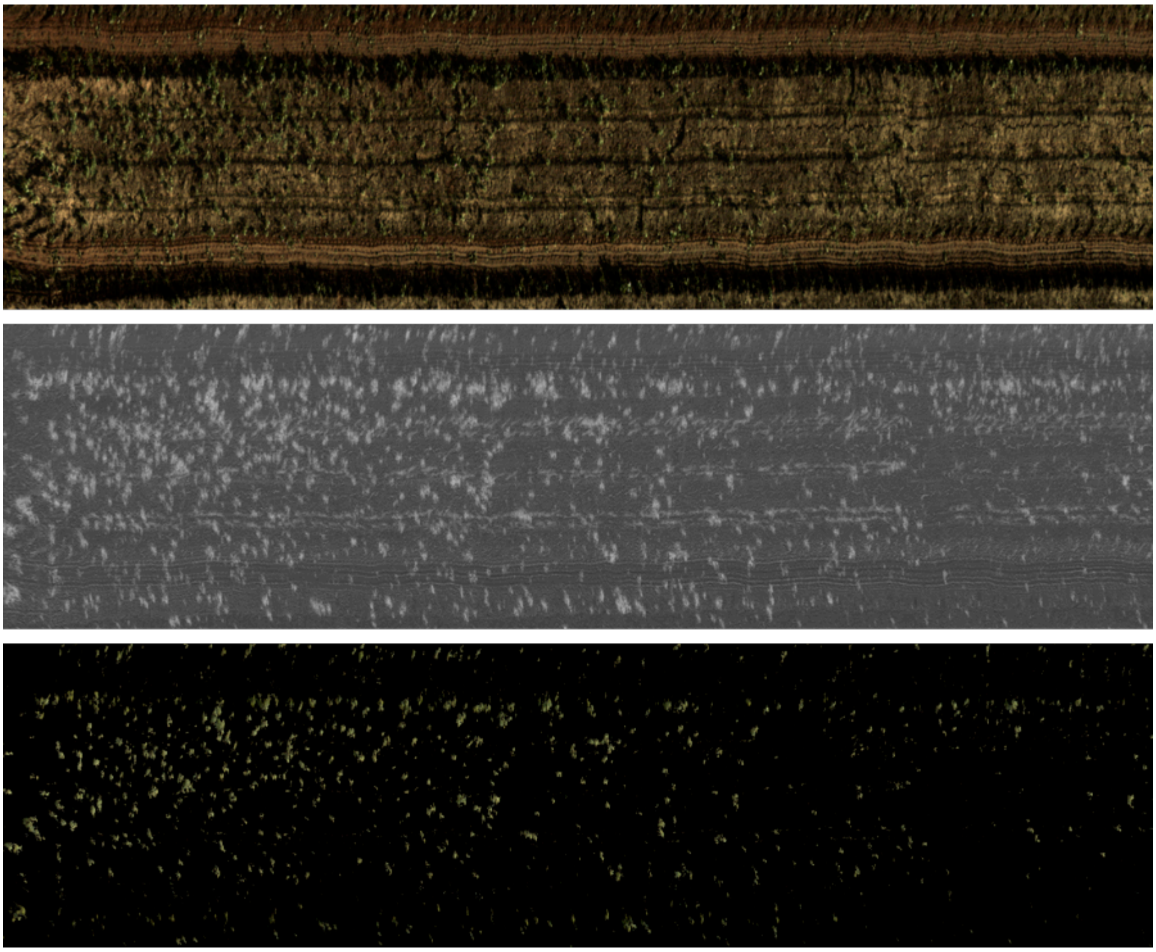

***Figure 18: Hyperspectral output and the corresponding NDVI and vegetation segmentation of the onion site.***

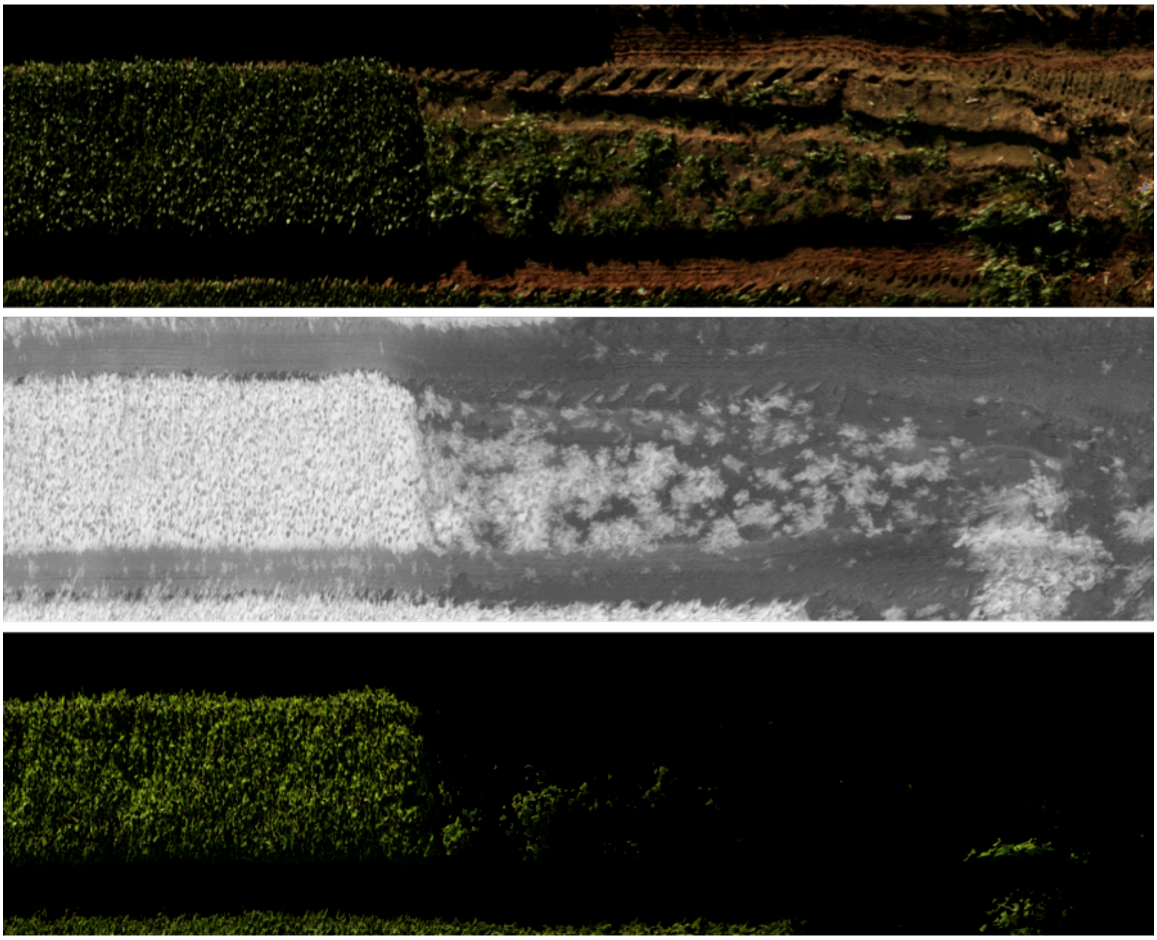

***Figure 19: Hyperspectral output and the corresponding NDVI and vegetation segmentation of the spinach site.***

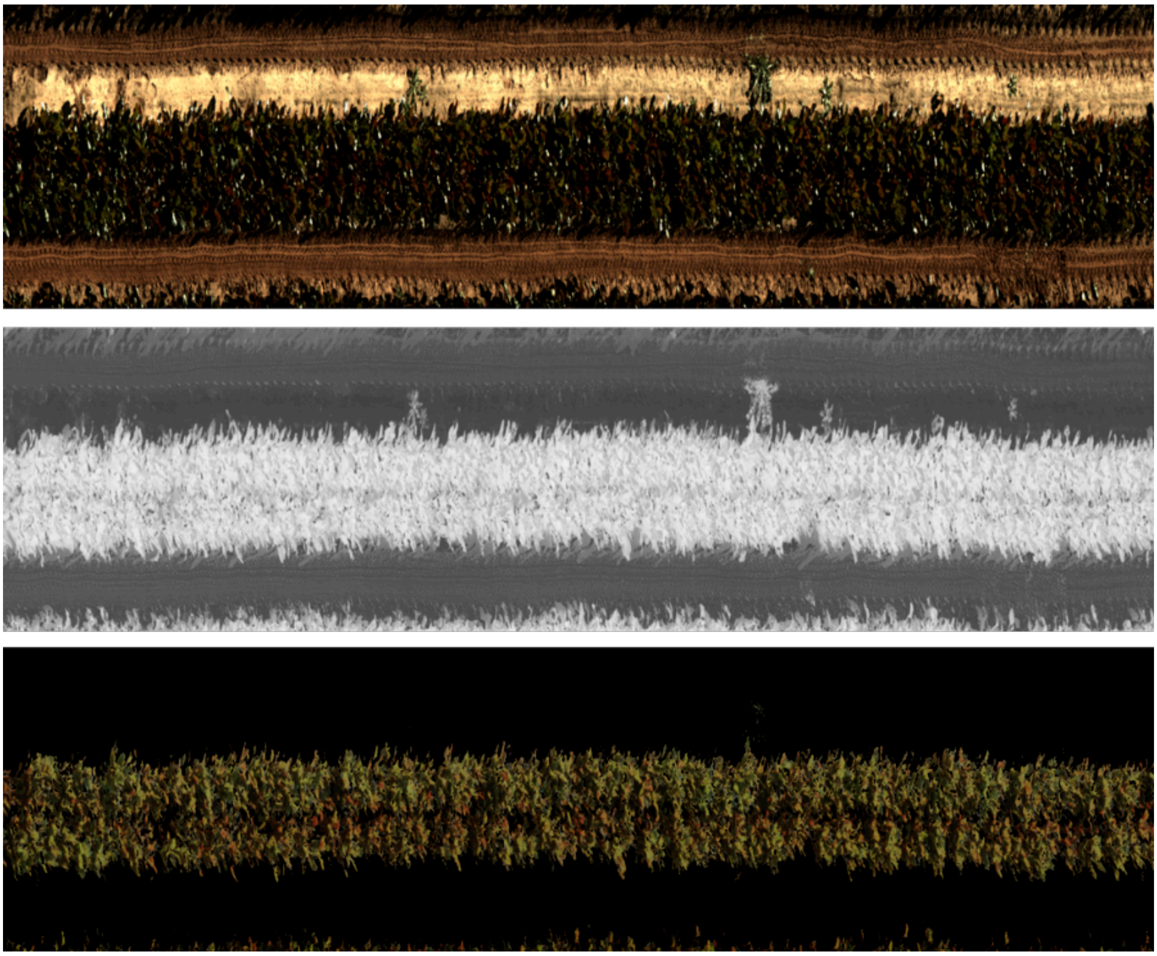

***Figure 20: Hyperspectral output and the corresponding NDVI and vegetation segmentation of the beetroot site.***

## Classification

This section outlines a vegetation classification pipeline, applied to onion and weed classification. Onions were chosen for this experiment because of the frequent occurrence of weeds observed on the four rows where herbicide had not been sprayed. By contrast, the surveyed areas of beetroot and spinach were relatively free of weeds.

Classification is tested on a section of the visual onion map. The image shown in Fig. 21 contains both weeds and onion leaves. Both are similar in colour (in visible and near-infra-red spectra), however, the weeds are larger leafier plants, whereas the onion leaves are taller, with only one or two 2-3mm thick leaves at this stage in the growing cycle.

The first step of the pipeline is to perform vegetation segmentation to extract the vegetation from the background soil. In this environment, the green colour (in the visible spectrum) was a reliable indication compared to the brown soil, thus standard RGB imagery was sufficient and the full hyperspectral sensing was not required for segmentation. The images are transformed from the original RGB colour space to hue-saturation-value (HSV), to improve robustness of the algorithm to changing illumination conditions. A band pass filter is then applied on the HSV image to segment the green regions. The resulting vegetation segmentation is shown in Fig. 22.

The next step is to perform a geometric analysis. Geometric properties including area, solidity and eccentricity are computed for each segment. The classifier then assigns class labels according to the geometric properties. The results are shown in Fig. 23. Vegetation that was classified by the algorithm as onion leaves are bounded by the blue boxes, whereas the weeds are bounded by the red boxes. The results show that the algorithm was able to accurately identify the weeds but was missing some of the onion leaves. This was due to the fact that the onions were still in an early growing stage and some of the plants were too small to be correctly identified, particularly when the leaves are oriented vertically towards the camera. The motion blur of the images also reduces the performance, because the resolution of the camera is effectively decreased (which was later addressed by a system of artificial illumination underneath the Ladybird).

*Figure 21: A patch of onion farm extracted from the onion map.*

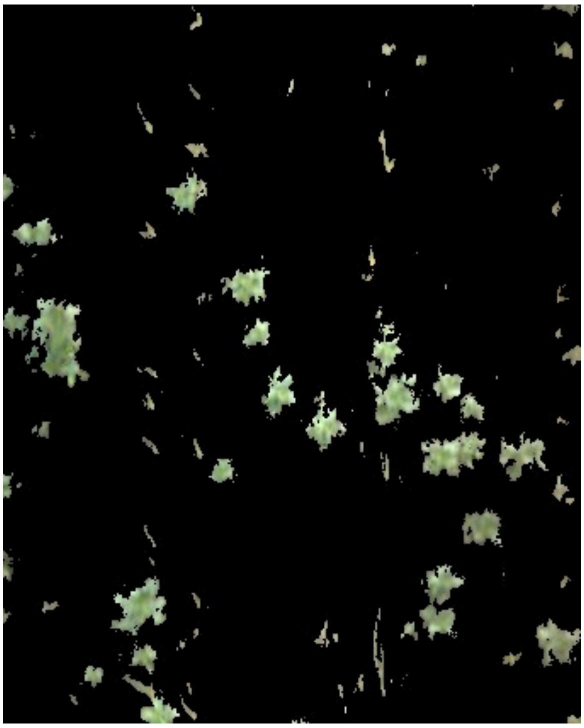

*Figure 22: Vegetation Segmentation*

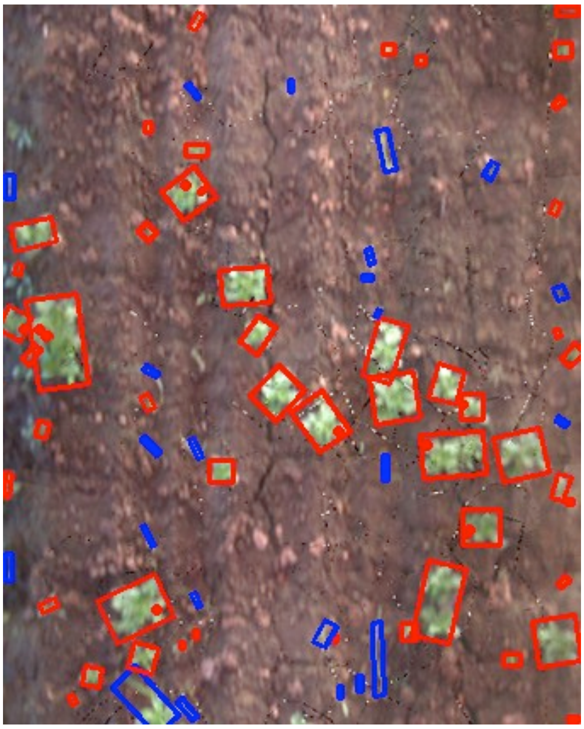

*Figure 23: Vegetation Classification. The onion plants are bounded by the blue boxes whereas the weeds are bounded by the red boxes.*

## Mapping the farm in 3D

Data from the forward and rear facing lidars can be combined with the localisation data from the Ladybird navigation system, to build 3D maps of the regions of the farm that were scanned by the Ladybird. Although the data in a single scan can be difficult to interpret, the complete 3D maps are clearer.

Fig. 24 shows the 3D point-cloud map for the section of 8 beetroot row-ends. Each point in this view (and all related figures) is a geo-referenced position, meaning the location on the Earth is known to within an accuracy of approximately 5cm. In the figure, the scene is rotated so that it can be viewed from above, and the points are coloured according to their local elevation.

Several phenomena are clearly observable from the data. In this area, the beetroot foliage is not uniform. For each bed (row), 3 sub-rows of beetroot seeds were originally planted, and although this is difficult to see from the foliage with the naked eye it is visible from the lidar data. Furthermore, there are two beds for which only two of the three sub-rows of beetroot are present, and within the rows with three out of three, there are some sparser patches of foliage. The 3D map is also shown for two full-length beetroot rows in Fig. 25, where it can be seen that a bed is missing one of the three sub-rows for the entire length of the field. This indicates the seed laying machine had a fault while laying seeds in this bed. Currently, growers manually count the vegetables per square metre in several randomly distributed patches, so the ability to measure this more accurately and more widely if of value, for monitoring and yield estimation. Although for root vegetables such as beetroot, where the yield is not directly observable, there is a relationship between plant vigour (the quantity of leafy foliage above ground) and the yield beneath the ground.

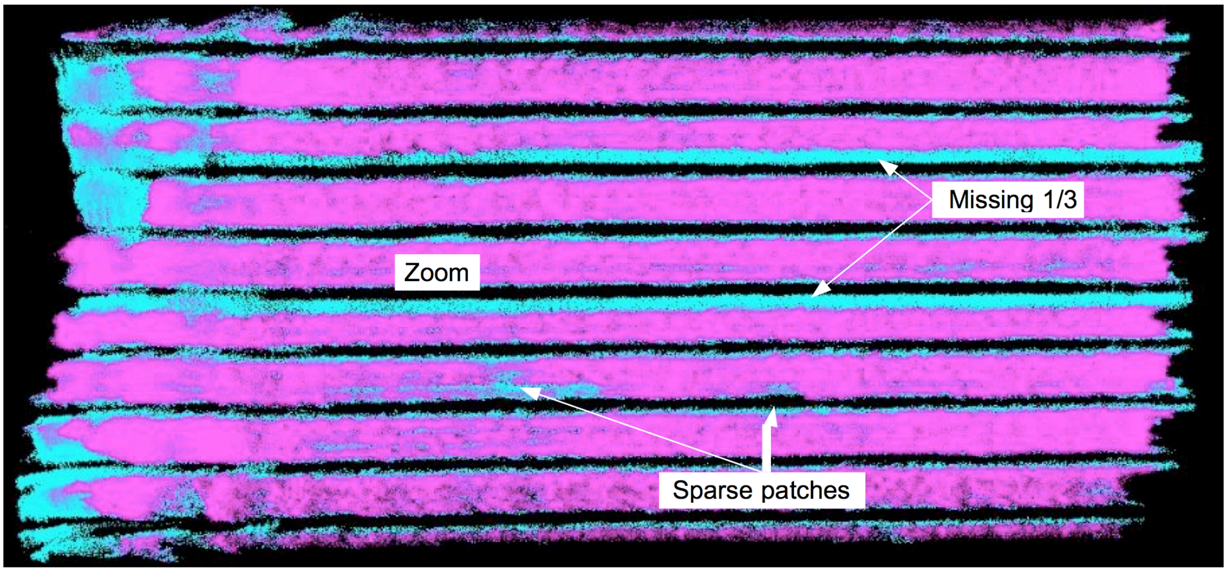


*(a) beetroot 8 sub-rows*

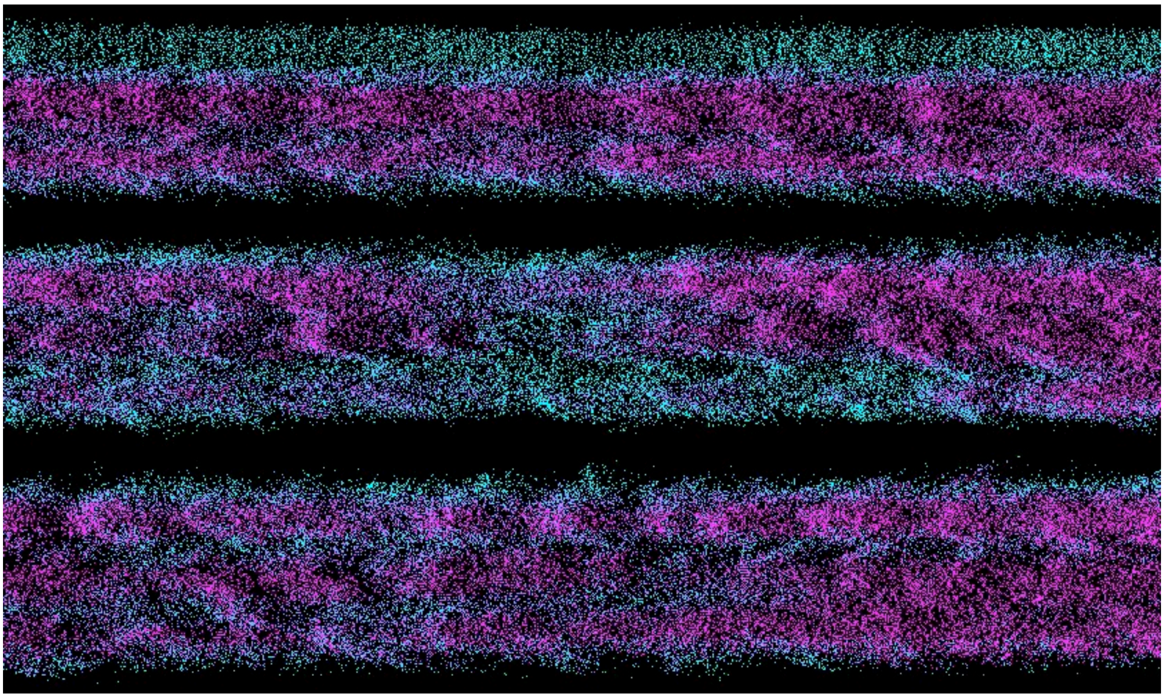

*(b) zoomed section*

*Figure 24: A 3D map of a section of beetroot, (a) wide and (b) zoomed. Geo-referenced lidar data is viewed from above. Points are coloured by elevation, from low (light blue) to high (pink).*

*(a) beetroot 2 full rows*

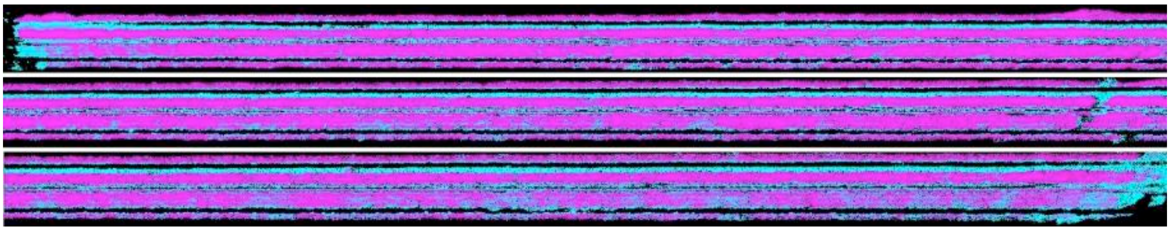

*(b) zoomed and split*

***Figure 25: A 3D map of two full beetroot rows, (a) wide and (b) zoomed and split to fit the page. Geo-referenced lidar data is viewed from above. Points are coloured by elevation, from low (light blue) to high (pink).***

Similar results are shown for spinach in Fig. 26. For this scan, there are no observed cases where an entire strip of vegetation is missing, as was observed for beetroot. A similar pattern of sparse patches is observed, however, it is possible that this is caused by a region of sunken ground, which may affect the geometry seen by the lidar, rather than indicating a reduced foliage. Unlike root vegetables, with bulk crops such as spinach, the yield is the leafy material above the surface, which is directly observable, therefore there is significant utility to using a system such as the Ladybird for measuring the yield in the field before harvest.

The 3D map for a section of onions is shown in Fig. 27. At the time of scanning, each onion had only one or two 2-3mm thick leaves, surrounded in some cases by sparse small weeds. This fine scale geometry is below the sensitivity of the lidar sensors, and therefore the map shows no variation due to foliage. There might be some utility for identifying regions of damage in the laser-levelled bed-structure, although the small sample region in the figure shows a uniform flat structure in this case.

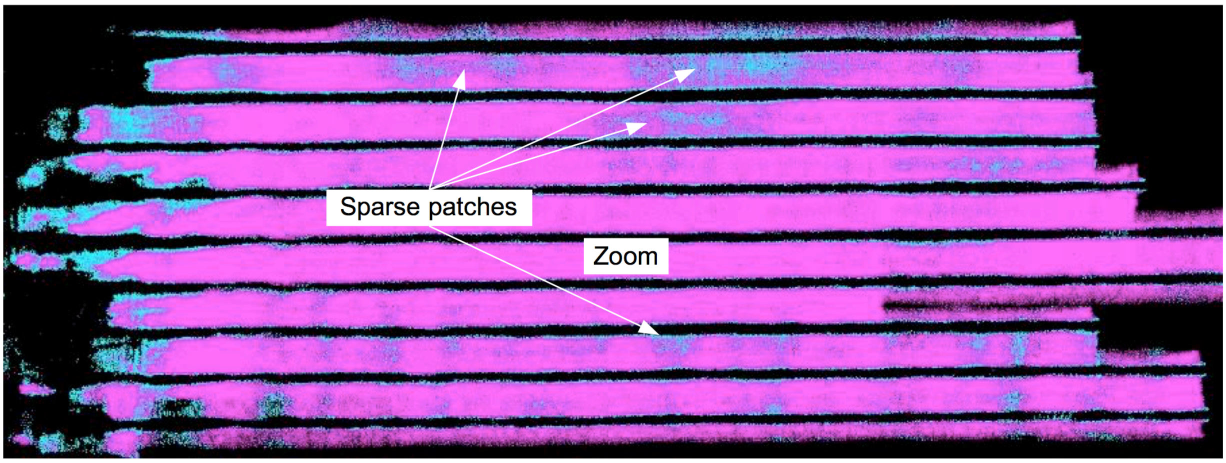


*a) spinach 8 sub-rows*

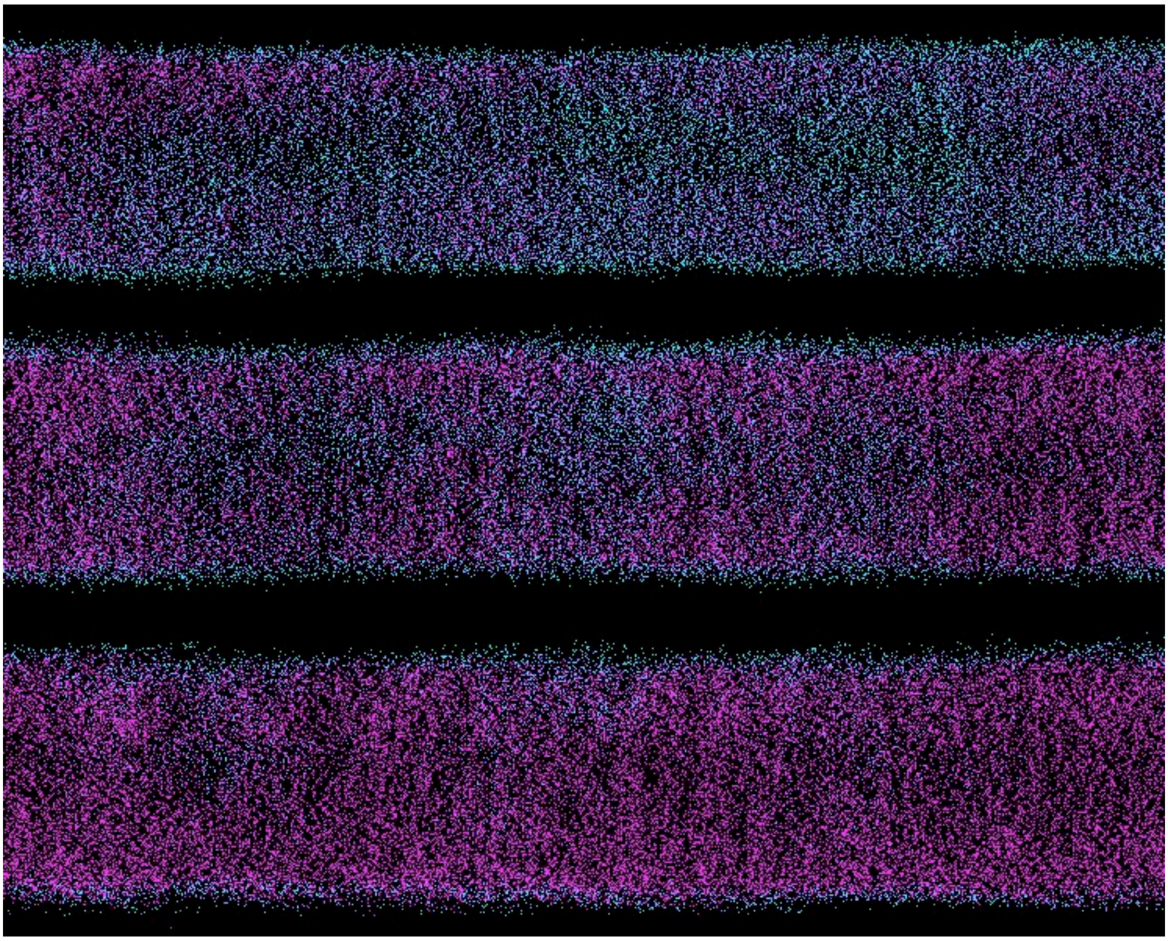

*(b) zoomed section*

*Figure 26: A 3D map of a section  of spinach, (a) wide and (b) zoomed.  Geo-referenced lidar data is viewed from above. Points are coloured by elevation, from low (light blue) to high (pink).*

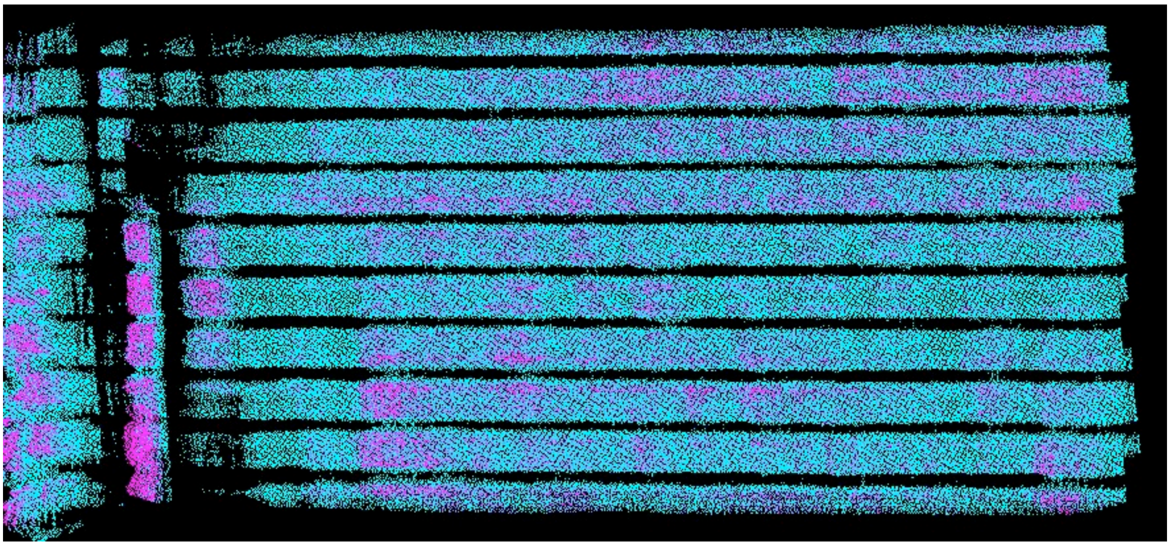

*Figure 27: A 3D map of a section of 8 rows of onions. Geo-referenced lidar data is viewed from above. Points are coloured by elevation, from low (light blue) to high (pink).*

The stereo visual cameras may also provide a useful, alternative source of 3D data for yield mapping applications. Fig. 28 shows an example of several subsequent stereo colour-range images, at a boundary between bare ground and the start of a row of spinach. Although there is noise in the data, the image suggests that this may be a feasible sensor for estimating the bulk volume of produce while it is still in the ground. It may be possible to calculate the volume estimates everywhere that the Ladybird surveys, to provide yield maps and totals.

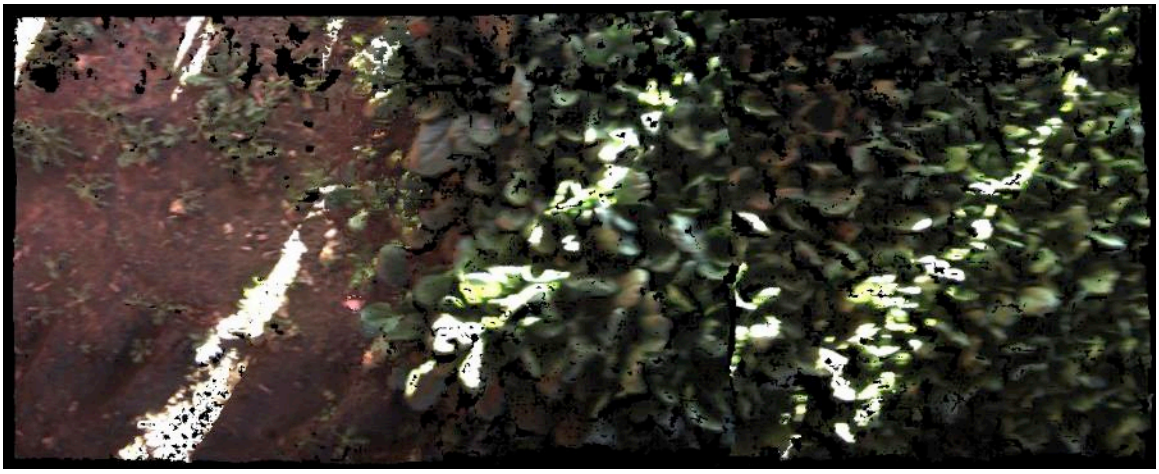

*(a) top view*

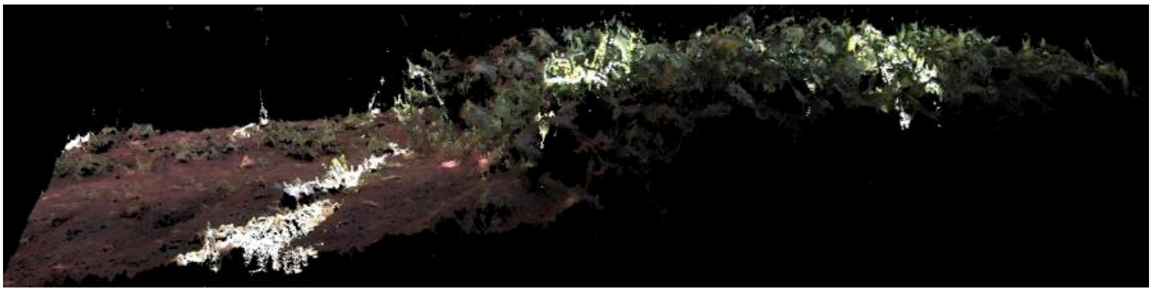

*(b) perspective view*

***Figure 28: Colour 3D stereo data, illustrating the potential for yield estimation for spinach. The scene from above (a) appears 'image-like' as this matches the perspective of the stereo camera. The 3D view is rotated in (b) to show the profile of the spinach. Volume esti- mates may provide a method to measure the yield of the crop while still in the field, as an alternative to lidar data.***

## *RIPPA*

With successful demonstration of the Ladybird Farm robot as a science platform demonstrating autonomous functionality, the next phase of the project moved towards developing a pre-production version of the robot. RIPPA™ (Robot for Intelligent Perception and Precision Application™) was developed as well as a the targeted shooting spray module proof of concept – VIIPA™ (Variable Injection Intelligent Precision Applicator™).

| ***Specification description*** | ***Value*** |
|---:|---|
| *Track width* | 1.52m |
| *Max crop height* | 0.6m (adjustable) |
| *IP rating* | IP65 |
| *Mass (no payload)* | Approx. 275kg |
| *Max payload* | 100 kg at max operating grade (12 deg) |
| *Charge-time from empty* | > 2 hours (dependant on charger) |
| *Idle discharge Time (no solar)* | 43 hours |
| *Driving discharge time (0.5m/s, no solar, no payload)* | 21.5 hours |
| *Max area traversed per charge (no solar)* | 8 hectares (~10 hours at 1.6m/s) |

### Lab Tests

Fig. 29 shows the first outdoor test of RIPPA on The University of Sydney's Darlington campus. The robot was joystick controlled by a human operator with the following key objectives:

- validate basic motion functionality
- demonstrate targeted spraying functionality

Basic motion functionality of RIPPA under joystick control was demonstrated by controlling RIPPA in a wide variety of its omni-directional motion capabilities. This early stage test provided a key validation point in the system development, with further enhancements to be made to the controllers and in transitioning towards an automated system.

On the same experiment, VIIPA demonstrated targeted spraying (Fig. 30). In the experiment, RIPPA is joystick controlled towards some weeds, and once stationary captures an image using an RGB-NIR camera. The centroids of vegetation targets (weeds) in the image are then located in image space, before the actuators control the nozzle in actuator space to apply a small dose of water to each of the targets. Fig. 31 shows an example NDVI image of vegetation (weeds) on ground (dirt/mulch) with fitted blue centre of mass ellipses, the centres of which are sequentially shot at by the nozzle. This same basic method is common to the rest of the experiments carried out through this report, with the addition however of added planning and perception intelligence to improve the performance and functionality of the VIIPA system.

The experiments conducted above demonstrated the targeted spraying robot with basic human operated motion functionality. The next stage of the project was to demonstrate the robot undertaking autonomous operation on a farm and spraying individual plants.

*Figure 29 - Initial demonstration of RIPPA (Robot for Intelligent Perception and Precision Application™) at The University of Sydney's Darlington campus.*

*Figure 30 - Underneath RIPPA™, showing VIIPA™ (Variable Injection Intelligent Precision Applicator™) in the top/middle of the figure.*

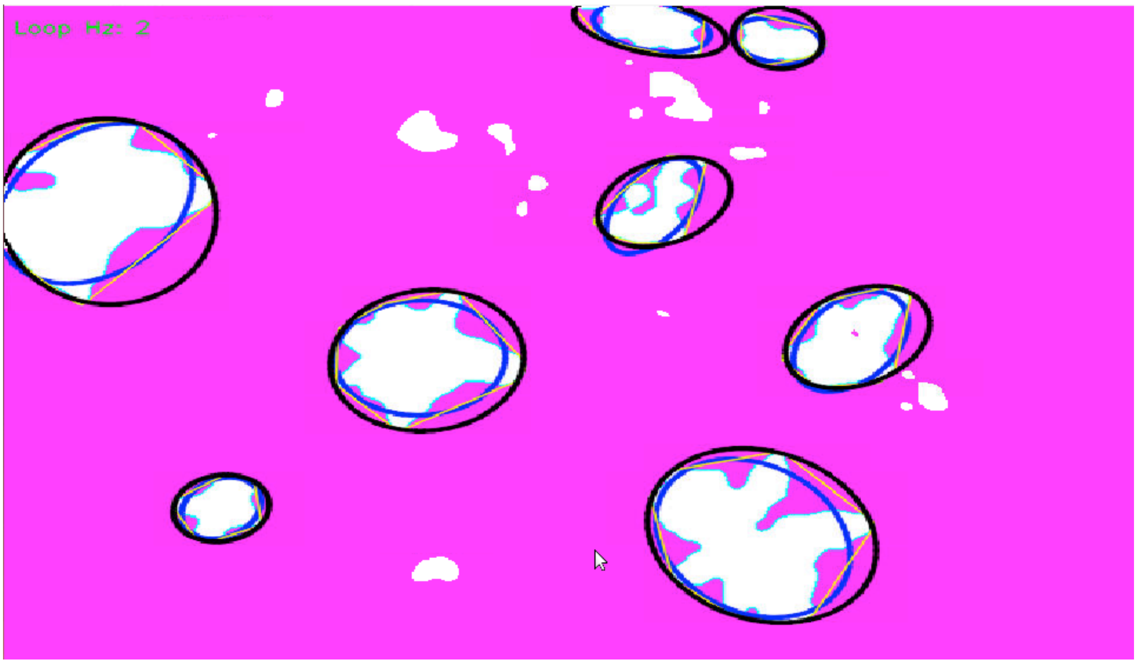


*Figure 31 - NDVI image of vegetation (weeds) on ground (dirt/mulch) with fitted blue ellipses representing the centre of mass and defining the target centre*

## Field Tests

RIPPA was demonstrated autonomously traversing a spinach crop in Cowra, NSW without the use of an RTK base station (Fig. 32). Two tests were conducted, the first at 0.4m/s and the second at 1.2m/s. The method of navigation for RIPPA for each of the experiments on the spinach crop was as follows:

1. Acquire accurate GPS solution to < 5cm range using Omnistar HP (support for Centerpoint RTX to be used in the future).
2. Map out the headlands by joystick operating along the southern headland and offsetting for the northern headland
3. Log the position of the starting row and interpolate consequent rows with a fixed track width offset (5ft)
4. Plan the path, in this case, the robot starts on theadland, slowly approaches the row, then traverses the row at target speed before slowing down ~30m into the row to return back to the start
5. Step iv) is repeated with the robot moving to each next row until the mission is complete

The experiment validated for the first time the autonomous navigation of the RIPPA robot using a GPS based correction service.


*Figure 32 - RIPPA™ on a spinach crop in Cowra, NSW in October 2015.*

Also on the spinach crop, the VIIPA module successfully demonstrated targeted micro-dose application onto the crop as shown in Fig. 33. For this experiment, RIPPA was joystick controlled onto the crop, and once stationary captures an image using an RGB-NIR camera. The centroids of vegetation targets in the image are then located in image space, before the actuators control the nozzle in actuator space to apply a small dose of water to each of the targets. This experiment uses NDVI to classify vegetation from not vegetation, though this alone is not capable of discriminating weeds vs crop.


*Figure 33 - RIPPA on the spinach crop in Cowra, NSW shooting vegetation using NDVI. This image is from a news article by Arabella Fingleton of WIN News Central West who filmed the event in Cowra, NSW on October 21st, 2015. [Online]. Available: https://www.facebook.com/431557963627074/videos/845773805538819/*

Figure 34 shows RIPPA on a corn crop in Cowra, NSW. The experimental setup for this test involved autonomously navigating RIPPA on the corn crop using RTK GPS via a fixed base station. The methodology of the experiment on the corn crop is as follows:

1. Acquire accurate GPS solution to < 1cm ε using a fixed RTK base station sending corrections over our own network.
2. Map out the headlands by joystick operating along the southern headland and offsetting for the northern headland
3. Log the position of the starting row and interpolate consequent rows with a fixed track width offset (5ft)
4. Plan the path, where RIPPA starts on the headland, slowly approaches the row, then traverses the row at 0.3m/s before slowing down ~30m into the row to return back to the start
5. While RIPPA is traversing the row, VIIPA is using the NDVI data as shown in Fig. 35 to classify and locate corn.
6. VIIPA then uses in built algorithms to aim and shoot the corn whilst moving as shown in Fig. 36.

This experiment validated for the first time the ability for VIIPA to automatically sense, predict, aim and shoot a crop using micro-doses of fluid. VIIPA demonstrated autonomously and accurately shooting vegetation (weeds and corn) with water in doses of 0.1g while moving at 0.3m/s. In a separate pre-trial test, VIIPA also demonstrated the ability to shoot pre-defined targets statically in doses as small as ~0.01g at up to ~13 targets per second with < 2mm precision and < 5mm accuracy.

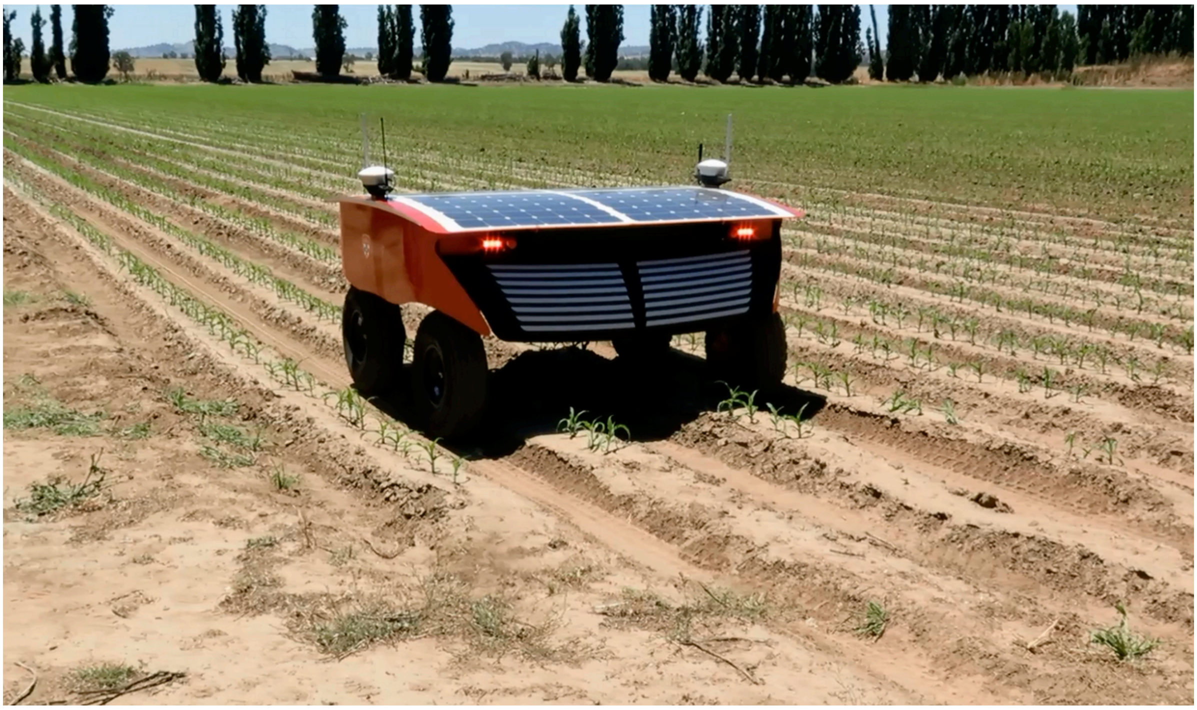

*Figure 34 - RIPPA™ on a corn crop in Cowra, NSW in November 2015.*

*Figure 35 - NDVI image with operations applied (e.g. size, morphological) to locate corn plants using RIPPA on a corn crop in Cowra, NSW in November 2015. Note here that the weed in the centre is found by the NDVI thresholding, though is filtered out as a crop, meaning it does not get shot at by VIIPA. In this image, only the 3 corn plants shown on the left wrapped with ellipses are shot with the target at the centre of each blue ellipse.*

*Figure 36 - VIIPA mounted to RIPPA and moving forward at 0.3m/s (in the upwards direction of the figure). The newly implemented intelligent perception system enables VIIPA to predictively apply a 0.1g dose of water to each corn while moving. The weed shown near the top centre of the image is not shot due to a size based filter.*

# Outputs

The key outputs from this project are:

- Demonstration of an intelligent robot for the vegetable industry, that is solar-electric and omni-directional.
- Demonstration of an initial set of algorithms for mapping, segmentation and classification.
- Demonstration of pre-production robot for the vegetable industry.
- Communications
  - PMA Australia + New Zealand, Keynote: Prof. Salah Sukkarieh http://www.pmafreshconnections.com.au/pages/program.php
  - ABC Rural, radio interview and online article "World first farm robot set to revolu- tionise vegetable farming" for ABC Rural, Alex Blutcher, 25th June 2014 http://www.abc.net.au/news/2014-06-25/farm-robotics-university-of-sydney-ve5550076
  - Panorama News Radio and Current Affairs (Melbourne), radio interview with Holly Thwaites, 25th June 2014
  - 2SCR (Sydney), radio interview, not yet aired.
  - The University of Sydney, news update, http://sydney.edu.au/news/84.html?newscategoryid=2&newsstoryid=13686
  - "Robotics at ACFR", routine updates and mailing list notification at http://sydney.edu.au/acfr/agriculture
  - BBC documentary and a live demonstration of the Ladybird was held at Mulyan Farms, Cowra NSW.
  - "Robo farmers - could they revolutionise agriculture?", BBC, http://www.bbc.com/news/technology-30143428. 25 November 2014 .
  - Live industry demonstration of the Ladybird was held at Mulyan Farms, Cowra NSW, for HAL and invitees from the industry. 8th September 2014.
  - WIN News Central West October 21st, 2015. [Online]. Available: https://www.facebook.com/431557963627074/videos/845773805538819/
  - Keynote, South Africa Hortgro Conference, Salah Sukkarieh, Smart Farms - Part 1, and Smart Farms - Part 2.
  - Vegetables Australia Magazine (HIA), "Experts gather for Summer School in ag robotics".
  - ABC Rural, 4th February 2015, Robotics to revolutionise farming and attract young people back to agriculture says Australian Centre for Field Robotics at Sydney University. http://www.abc.net.au/news/2015-02-04/agricultural-robotics-future-jobs/6068450
  - The Australian Centre for Field Robotics (ACFR) at The University of Sydney hosted the inaugural IEEE RAS Summer School on Agricultural Robotics, 2015. Organisers: Fitch, R. (ACFR), Sukkarieh, S. (ACFR), Bergerman, M. (CMU), van Henten, E. (Wageningen University).

# Outcomes

The key outcomes from this project are:

- Design, development and demonstration of an intelligent robot for the vegetable industry, that is solar-electric and omni-directional.
- Design, development and demonstration of an initial set of algorithms for mapping, segmentation and classification. These algorithms were able to distinguish individual crops, segment out the ground from the crop, and classify crop from weeds, including the ability:
    - To autonomously gather and process data over different vegetable crops,
    - To detect and map seedlings and calculate their statistics,
    - To detect and map vegetation and to measure the abundance of leafy crops to provide yield distribution maps.
- Design, development and demonstration of pre-production robot for the vegetable industry that was able to distinguish weeds from crops and accurately spray the weed:
    - Including the functionality of joystick and autonomous operation and automatic detection of weeds using NDVI imaging,
    - Demonstrations of the robot undertaking continuous operation on the farm spraying individual plants - weeds and crop - of different varieties.

# Evaluation and Discussion

This report has presented the background, objectives, methodology, outputs and outcomes for project VG12104 *An intelligent farm robot for the vegetable industry*.

The report detailed the design, development and demonstration of the Ladybird Farm Robot. The system was designed with modularity at heart and this has proven to be an extremely effective principle. The robot can be reconfigured easily to operate on different farm row widths across different crops. The system developed also aimed at testing the effectiveness of solar-electric principles for farm operations, again demonstrating that such technology is highly effective. The various sensors on board provide a significant amount of multi-modal data from which various data analysis and machine learning can be conducted. This has led to the development of algorithms for mapping, segmentation and classification. Although the results have shown the potential for such technology, further study is required to ground truth the algorithms in order to measure their level of accuracy, and to determine the payoff between accuracy and computation effort / real-time performance.

The report has also detailed the newly developed commercial prototype RIPPA including the demonstrations of RIPPA with VIIPA on a spinach and corn crops. RIPPA was demonstrated driving autonomously, along with intelligent predictive spraying while moving. VIIPA demonstrated autonomously shooting vegetation (weeds and corn) with water in doses of 0.1g while moving at 0.3m/s. In other experiments, the VIIPA system also demonstrated the ability to shoot pre-defined targets in doses as small as ~0.01g at up to ~13 targets per second with < 2mm precision and < 5mm accuracy. The preliminary results of the above experiments, show promise for future applications in using RIPPA and VIIPA within vegetable crops for task automation.

Feedback for the architecture of both robotic systems has been sought by speaking with key Australian growers (e.g. visiting farms and field days) where two-way communication between engineers/researchers and growers helps both ready robots for farms, and farms for robots.

The learning from this project to date is that there exists significant commercial and research based opportunities for both RIPPA and VIIPA in the vegetable industry and Ladybird for the scientific/plant industry.

# Recommendations

This project has demonstrated the capability for autonomous robots to acquire and utilise information that has been gathered in real-time on the farm. The systems developed have been robust and the initial algorithms have demonstrated the potential for intelligent decision making in real time.

It is recommended that further research and development into automated solutions that would capitalise on systems developed in this project should be considered.

In particular the aim should be to research, develop, evaluate and support the adoption and commercialisation of automation technologies that aim to reduce production cost and increase on farm productivity in the vegetable industry. These technologies will collectively provide timely and accurate information about crop status (including yield, quality and forecasting) and minimise input costs through the precision application of chemicals and water (and in some cases the use of non-chemical automated systems for precision crop and weed handling).

The recommendation would be to divide the project into four themes:

- Sensing – further study into sensors for soil, water and nitrate mapping (hyperspectral, VNIR, SWIR), including potential use of thermal imaging (LWIR) and fusion with other modalities such as stereo vision; and ground penetrating radar sensors to measure in-ground traits including soil water content and sub-surface plant traits (e.g. the root system).

- Automated Decision Support Systems – further study into better yield estimation models over a wider variety of crops and multiple years; crop quality mapping and automated decision models to achieve crop uniformity; automated crop forecasting to predict optimum harvest time, or selective harvest where appropriate.

- Crop Interaction – conduct extensive demonstrations of the RIPPA with the VIIPA system at various grower and adoption sites to demonstrate the capability of autonomously targeting crops and dispensing fluid accurately.

- Farm Automation Standards – In order to successfully adopt new technologies into vegetable production systems and to minimise silo solutions (i.e. to allow plug-and-play of various technologies to a working system), horticulture technology standards need to be developed, and where required identification and development of policies to meet regulations.

.

# Scientific Refereed Publications

Conference article and workshops

- Underwood, J.P., Calleija, M., Taylor, Z, Hung, C., Nieto, J., Fitch, R., Sukkarieh, S., 2015. Real-time target detection and steerable spray for vegetable crops. Workshop on Robotics in Agriculture at the International Conference on Robotics and Automation (ICRA).
- http://www.seas.upenn.edu/~tokekar/ICRA2015Workshop/
- Fitch, R., Sukkarieh, S., Bergerman, M., van Henten, E. "2015 IEEE RAS Summer School on Agricultural Robotics (SSAR2015)". http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7124616

# Intellectual Property/Commercialisation

The targeted shooting / spray system VIIPA described throughout this report is a specific implementation of Australian patent #2014905005 filed 10th December, 2014 by The University of Sydney.

Hardware and software artefacts for the design of the robots and the algorithms for automated operation, mapping, classification and segmentation, are under licensing agreements with the University of Sydney and HIA.

# Acknowledgements

This work is supported by the Australian Centre for Field Robotics at the University of Sydney, AUSVEG and Horticulture Innovation Australia through project VG12014 An intelligent farm robot for the vegetable industry. The authors would like to acknowledge the substantial contributions of Vinny Do, Vsevolod Vlaskine, Andrey Sokolov and Navid Pirmarzdashti for their work in software development, Bruce Crundwell, Victor Chan and Ben Stewart for their assistance in hardware development, Steven Potiris and William Reid for their contribution to the development of VIIPA and field trial support.