final report

Project code: B.NBP.0474

Prepared by: Mitch Bryson and Salah Sukkarieh

Australian Centre for Field Robotics, University of Sydney

Date published: October 2011

ISBN: 9781741916447

PUBLISHED BY
Meat & Livestock Australia Limited
Locked Bag 991
NORTH SYDNEY NSW 2059

UAV Surveillance Systems for the Management of Woody Weeds

Meat & Livestock Australia acknowledges the matching funds provided by the Australian Government to support the research and development detailed in this publication.

Abstract

Woody weed infestations that cover large open grassland areas are difficult to control and their management is extremely costly and time consuming. The aim of this project was to develop and test Unmanned Aerial Vehicle (UAV) surveillance systems which can detect woody weed infestations and provide the information needed for woody weed control and eradication. The project focussed on the sensor and data fusion algorithms needed for weed detection and the precise positioning of a dispenser for the distribution of herbicide at the infestation point. Two types of UAV systems were used in the project, a fixed-wing UAV (FUAV) for broad acre surveillance, and a rotary-wing UAV (RUAV) for precision targeting around the weed. Results of the project demonstrated the ability to build high-resolution maps over farmland areas and classify different species of weeds within the maps using vision data collected from the UAV, and the ability to hover of specific geo-referenced locations of where weeds were detected for the distribution of the herbicide.

Executive summary

Woody weed infestations that cover large open grassland areas are difficult to control and their management is extremely costly and time consuming. The aim of this project was to study the use of surveillance systems on Unmanned Air Vehicles (UAVs) to provide the potential for an all round effective and reliable means to manage and eradicate infestations.

The objectives of the project were to develop and test various image recognition algorithms that could detect and geo-locate woody weed infestations in open farmland and to demonstrate the algorithm's potential by using a fixed-wing UAV (FUAV). Furthermore, the project aimed to test a novel herbicide distribution system that could be mounted on a rotary-wing UAV (RUAV), which would hover over the geo-located weeds.

The project ran over three years (2008-2010). The first year focussed on the development of the surveillance system and determining the location of the demonstration sites. Two sites were selected: Carrum Farm and Williams Outstation, both located within the vicinity of Julia Creek, QLD.

In 2009 the first demonstration was conducted over these two properties. The aim was to firstly establish that the UAV systems would operate effectively in this environment, and secondly to collect data in order to research into various classification and identification algorithms. A sensor payload consisting of a monocular colour vision camera, Inertial Measuring Unit (IMU) and Global Positioning System (GPS) receiver attached to the FUAV was used to collect data in these areas to test the algorithms. The algorithms developed showed that the system was capable of building 3D terrain images with a spatial resolution of 3cm per pixel and 1-2m accuracy (with respect to GPS reference coordinates) and to discriminate accurately between different species of vegetation, including weed species versus native tree species using the imagery and map data.

Project work in 2010 focussed on preparations for the final experimental demonstration of the UAV systems for weed surveillance and herbicide delivery at Julia Creek. Based on the prototype mapping and weed classification algorithms developed in 2009, the aim of the 2010 flight trials was to demonstrate that these algorithms could potentially build classified weed maps with an overnight turn around time (within 24 hours) after data collection. Work thus focussed on improving the accuracy and robustness of the algorithms and to produce software implementations that required minimal human input to the mapping and classification process. The FUAV used in the 2009 flight trials was once again used in the final system demonstration to collect data for map building and weed classification, with no changes made to the sensor payload. During the final demonstration classified weed maps were built within 24 to 48 hours after data collection, demonstrating that overnight processing with minimal human input is feasible. Work also focussed on building and flight testing the RUAV for granular herbicide delivery. A computer based autopilot system including GPS for guiding the RUAV to a fixed location was integrated into the system. The aim of the final demonstration was to show the RUAV navigating to fixed locations of weeds provided by the weed map, autonomously built using the FUAV data and to show the granule dispenser in action. The RUAV was demonstrated flying autonomously, and the herbicide delivery system was demonstrated separately.

Results from the project demonstrated that several important weed species such as Prickly Acacia, Parkinsonia and Mesquite could be effectively distinguished and mapped from native species. Future work will focus on further accuracy discrimination between weed species themselves and on using the rich 3D terrain data collected from the platform for a variety of other farming and agricultural applications.

Contents

1 Background

Woody weed infestations that cover large open grassland areas are difficult to control and their management is extremely costly and time consuming. This project has studied the use of surveillance systems on Unmanned Air Vehicles (UAVs) to provide the potential for an all round effective and reliable means to manage and eradicate infestations.

2 Project objectives

1. Demonstration of real time image recognition algorithms and geo-location of woody weeds from a fixed-wing UAV in a 5x5km test area.
2. Demonstration of a hovering UAV manoeuvring to geo-located weed points and precise positioning within the 5x5km test area.
3. Demonstration of weed control through pinpoint placement of herbicide at weed locations
4. Recommendations for future research

3 Methodology

3.1 Field Trip to the Mitchell Grasslands, Queensland

A field trip to Cloncurry and Julia Creek was conducted in August 2008. The purpose of the trip was to allow ACFR staff (Salah Sukkarieh) to gain first hand knowledge of the environment, weed ecology, and management processes of woody weed infestations. A number of sites were chosen by Nathan March (National Coordinator for Mesquite, Parkinsonia and Prickly Acacia) who also coordinated the meetings with the various farm managers. Rodd Dyer (MLA project manager) was also present. Among the various sites visited, including Fort Constantine, Dalgonally, and the DPI&F reserve, two sites, Carrum Farm and Williams Outstation, were chosen as ideal as they contained the three weeds of Mesquite, Parkinsonia and Prickly Acacia amongst them, and had relative easy terrain and open space so that the UAV experiments could be conducted. One important point that came from the trip was the timing of the flight trials. The opportune window for conducting the tests is during the dry season, and specifically from May to August. This provided the best discriminatory response from any data collected and given that it was also the mildest time of the year would make it easy to conduct the tests and get support from farm ground staff.

3.2 Preparation of Flight Platforms and Sensor Payload for Weed Finding

The first half of 2009 saw an extensive period of preparation for the flight platforms involved in the project, focusing on the data collection requirements. Effort this year was focussed primarily on preparation of the fixed-wing UAV system (the J3 Cub, see Figure A1) with full integration of the necessary sensor payload and engineering of the platform for operations on an improvised runway. Preparation of the fixed-wing UAV system was driven by the following requirements for both mapping and classification of weeds:

1. To build 3D spatial maps of weed infestations and other objects in the terrain over approximately a 5x5km area with precise positioning (a spatial accuracy of 5m was chosen).
2. To collect imagery over the area at a spatial resolution of at least ¼ of the smallest weed patch (approximately 25cm per pixel).

Preparation of the fixed-wing UAV system included integration of a sensor payload box consisting of an Inertial Measuring Unit (IMU), Global Positioning System (GPS) receiver and a downwards-mounted colour vision camera system (see Figure A2 and Table A1 for sensor specifications). The combination of IMU and GPS allowed for the estimation of the own-location of the UAV (i.e. its own position and orientation along the entire flight trajectory). This information is critical in referencing the position in the terrain of objects (such as weeds) identified in the image coordinates. Additionally it was decided to integrate information from the camera into the estimation of the own-location of the UAV in order to improve mapping accuracy results (a technique known as Simultaneous Localisation And Mapping (SLAM)); this introduced a requirement that the vision camera frame captured would be at least at 50% of the subsequent video frame.

The area coverage required and the spatial resolution of the imagery was the main driver for the resolution of the colour camera system and the endurance and operating altitude of the flight platform. Two different operating altitudes were chosen in order to survey differences in coverage and spatial resolution; the first operating altitude being 100m above ground-level and the second being 500m above ground-level. Based on the camera resolution of 1024x768 pixels and the use of a 12.5mm camera lens with a field of view of approximately 28x22^{o}, this provided an expected spatial resolution of approximately 3.7cm/pixel and a ground footprint of 38m at 100m altitude and spatial resolution of approximately 18.6cm/pixel and a ground footprint of 190m at 500m altitude (see Table A1). Primarily data was used from flights at the lowest altitude and thus highest resolution imagery. These values were significantly better than the original requirements specified, but seeing as they seemed achievable at very little detriment to the other requirements and that the applicability of the vision-based classification was the least certain aspect of the system, it was decided to aim for these values.

Based on the weight of the payload and onboard fuel, it was found that a flight time of 60 minutes at a speed of 25m/s was appropriate based on the data collection requirements. Based on the calculated footprints, a flight pattern of successive swaths up and down the flight area with approximately 20% lateral overlap was performed in each flight. It was found that in one flight the system could cover an area of 4000x600m at 100m altitude and an area of 6000x1500m at 500m. Covering a 6000x1500m area at each site allowed us to meet the requirement for a total mapping area of approximately 24 square kilometres (approximately equivalent to a 5x5km area).

Additionally to the onboard sensors, the UAV also carried two PC104 computer processing stacks, the first stack was used to log data from the IMU and GPS sensors whereas the second stack was used to log information from the colour vision camera. Although the takeoff and landing of the aircraft was performed by a pilot using remote control, flight control of the vehicle for following the fixed paths specified in the flight plans was performed using a commercial of-the-shelf autopilot system, the cloud cap piccolo, which allowed for autonomous operations beyond line-of-sight of the pilot and ground station crew. Once each flight was finished, data was downloaded to laptops at the ground station for processing.

3.3 July 2009 Flight Trials at Julia Creek, Queensland

During July, field trials to Julia Creek area were undertaken by a group from the ACFR comprising of three technical staff (Muhammad Esa Attia, Steve Keep and Jeremy Randle) two research staff (Mitch Bryson and Ali Goktogan), four research students (Calvin Hung, Nick Lawrence, Alistair Reid and Zhe Xu) and the project leader (Salah Sukkarieh). Two different flight sites were designated, one at Carrum farm and the other at Williams Outstation, where a variety of different flight plans were performed at each site. All of the ACFR team were involved in various activities such as managing/operating the platform, operating the ground station during flight, managing/operating the sensor payload box, processing data after flights and performing ground-based manned surveys of different species of weeds in the flight area as a source of ground truth and comparison data. Eight flights were performed over a period of 3 days, each flight lasting approximately one hour (Table 1).

Flight Number	Date	Area	Altitude
Flight 10	17^{th} July	Carrum	100m
Flight 11	17^{th} July	Carrum	100m
Flight 12	18^{th} July	Williams Outstation	100m
Flight 13	18^{th} July	Williams Outstation	100m
Flight 14	18^{th} July	Williams Outstation	500m
Flight 15	19^{th} July	Carrum	500m
Flight 16	19^{th} July	Carrum	100m
Flight 17	19^{th} July	Carrum	100m

Table 1 – Summary of Flight Operations during the July 2009 Flight Trials

In addition to flight operations, the field trials were also used as an opportunity to perform a ground-based survey of different types of woody weeds and other native vegetation under the flight paths, data which could be used to both train and test the classification algorithms developed. Members from the ACFR staff went out by four-wheel-drive to the flight area with weeds coordinators (Nathan March and Claire Dyason) and were trained in how to identify different types of woody weeds such as Prickly Acacia, Mesquite and Parkinsonia. Based on this knowledge, ACFR staff performed several ground-based surveys in different parts of the flight area. A handheld GPS was used to survey the position of several shrubs, where the type of shrub was recorded along with a photograph, which could be used to verify the classification made (see Figure A9).

Figures A10 and A11 illustrate maps of the ground truth data collected over both the Carrum farm and Williams outstation sites. At both locations a large density of the weed Prickly Acacia was observed and thus the ground survey focussed on labelling all other surrounding vegetation which was not Prickly Acacia (i.e. Parkinsonia and other native trees and plants such as Eucalyptus trees). The collected survey data was later used to help train the classification algorithms; the collected survey data was overlaid with the map imagery collected by the UAV and used to identify examples of different weeds from the aerial footage. This small collection of examples was then used to train the classification algorithms to detect weeds across the entire mapped area.

3.4 Preparation of Flight Platforms and Sensor Payload for 2010 Trials

The J3 Cub fixed-wing UAV (FUAV) system (see Figure A1, A2) used in the 2009 flight trials remained unchanged in 2010 and thus the majority of the work focussed on developing the rotary-wing UAV (RUAV) system (see Figure A3) for herbicide delivery. The RUAV was fitted with an off-the-shelf autopilot system including IMU and GPS for flight control. Additionally it was also fitted with a secondary GPS receiver which was differentially corrected using the same DGPS base station as the fixed-wing UAV, in order to provide sub-meter positioning accuracy required for navigating to geo-referenced weed locations. A remotely controlled granule dispenser was designed, built and mounted to the nose of the hovering UAV (see Figure A4). The dispenser was controlled by a servo which was linked through the autopilot ground control station such that granules could be dispensed automatically when the UAV hovered over a specified waypoint. The dispenser allowed for differing quantities per delivery and could hold a total granule mass of 400g. At the recommended dosage of 2.3 grams for an average sized tree, this allowed for delivery of approximately 175 doses before refilling the dispenser. In practice, the number of trees that could be treated in a single flight would depend on the density of infestation and the amount of time required in moving from one tree to another.

3.5 August 2010 Flight Trials at Julia Creek, Queensland

During August 2010, field trials to Julia Creek area were undertaken by a group from the ACFR comprising of three technical staff (Muhammad Esa Attia, Steve Keep and Jeremy Randle) one research staff (Mitch Bryson), four research students (Calvin Hung, Alistair Reid, Zhe Xu and Prasad Hemakumara). FUAV flights were made at each of the two flight sites designated during the 2009 trials with two flights at Carrum farm and three flights at Williams Outstation. All FUAV flights were flown at a nominal altitude of 100m above the ground. Additionally, the RUAV was flown at Carrum farm (see Table 2 and Section A.4 for details on flight paths). As in 2009, a ground survey at each flight location was also performed in order to collect reference data on the different types of weeds in the flight area. The location of different weeds and the weed type was recorded using handheld GPS receivers. This ground truth data could therefore be located in the flight imagery based on the geo-referenced map and used to assist in classifier training and for verification of classification accuracy.

UAV	Flight Number	Date	Area	Notes
Fixed-Wing	Flight 27	01/08/2010	Williams Outstation	
Fixed-Wing	Flight 28	01/08/2010	Williams Outstation	HDD Logging issue: data not used
Fixed-Wing	Flight 29	02/08/2010	Carrum Farm	
Fixed-Wing	Flight 30	02/08/2010	Carrum Farm	
Fixed-Wing	Flight 31	03/08/2010	Williams Outstation	
Hovering	Flight 01	05/08/2010	Carrum Farm	

Table 2 – Summary of Flight Operations during the August 2010 Flight Trials

The flight trials were also used as a chance to demonstrate the turn around time required to produce geo-referenced maps with classified weeds using the flight data. Data collected during Flight 27 at Williams Outstation was processed on-site using two laptops computers and three operators (Mitch Bryson, Calvin Hung and Alistair Reid) operating prototype software developed during the first half of 2010. Within 24 hours a geo-referenced map of the flight area was constructed and within 24-48 hours this map was populated with detected and classified weeds. Data from the remaining flights was processed partly while on flight trials and partly at the ACFR in the weeks following the trials.

In addition to the fixed-wing flights, the RUAV system was flown at Carrum farm on the final day of the flight trials.

3.6 Development of Weed Mapping and Classification Algorithms

Development of algorithms for detecting weed infestations began in 2009 after data collection from the 2009 flight trials and proceeded in two areas. The first area focused on using the collected sensor data to build an accurate geo-referenced map of the environment on which vision data could be overlaid, where the second area focused on classifying data in the vision frames into different types of vegetation including different types of woody weeds. After prototype algorithms were developed in 2009, the main aim for mapping and classification algorithm development during 2010 was to improve the accuracy, robustness and processing time requirements such that in-field operation was possible. These algorithms were implemented into software suitable for running on off-the-shelf laptop computers for data processing that could potentially build maps of classified vegetation within 24 hours of data collection. The pipeline for building a weed map consists of the following high-level steps:

(1) FUAV Data Collection: The UAV is flown over a target area carrying a sensor payload consisting of a colour vision camera, IMU and GPS. Sensor data is logged to a hard drive mounted on the UAV during the flight.

(2) Data Download and Distribution: Once the UAV has landed, logged sensor data is copied to multiple portable hard drives for processing and backup.
(3) Imagery Geo-referencing and Map Building: All sensor data is used to build a geo-referenced 3D terrain map corresponding to objects seen in the image data.
(4) Image-Based Tree Crown Detection and Classification: Image processing techniques are firstly used to identify the locations of tree crowns in the image data. A supervised learning process is then used to classify identified tree crowns into one of a number of species classes specified by a human user.
(5) Final Classified Map Delivery: Once tree crown classifications in the imagery have been made, they are transferred to the geo-referenced map which displays to the user the location and distribution of the identified weeds in a global coordinate system.

The overall data processing pipeline is shown in Figure A5.

3.6.1 Data Collection and Distribution

Once the UAV landed after a given mission scenario, data was distributed to a number of ground processing laptops via portable hard drives. Due to the large amounts of high-resolution imagery collected, the time spent copying data became an important issue in the speed of data processing. Therefore a system was developed where IMU and GPS data were copied first (due to their relatively smaller file sizes), where processing began on one of the laptops in parallel to image data being copied via a second laptop. Image data was then distributed to two laptops for parallel image processing and classification.

3.6.2 Imagery Geo-referencing and Map Building

The first stage in map building is to process IMU and GPS receiver data to compute an initial estimate of the pose trajectory of the platform during flight. This is performed using an Extended Kalman Filter (EKF) to estimate the position and orientation of the UAV at the time of each camera measurement. In parallel to navigation sensor processing, information from the camera is processed for mapping. SIFT image features were extracted from each of the images captured in flight and matched between subsequent frames. Once initial pose estimates from the navigation data have been computed, they are used to generate a list of potentially overlapping image frames, based on the estimated ground height of the UAV. SIFT features were then matched across multiple frames in the image data captured at different sections of the flight (referred to as “loop closures”). Once image matches have been computed for all the image data, an optimisation/estimation procedure was used to compute an optimal estimate of the UAV poses using features matches and the IMU/GPS data. These optimal pose estimates are then used to create a photo mosaic of the flight area which can be used to geo-reference an object detected in the image data. Additionally, a 3D model of the terrain can be generated by triangulating a set of dense feature points between subsequent frames based on the optimal pose data.

3.6.3 Image-Based Tree Crown Detection and Classification

In parallel to map building, separate laptop computers were used to perform tree crown detection and classification using the collected image data. We employed a tree crown detection method, which segments images into regions of background and trees. Individual tree crowns were then detected using an object-based template method which exploits information about the projected shadow of the tree (see Figure A6). The locations of tree crowns were then passed to a second classification phase. The classifier was based on a supervised learning approach in which human-identified examples of different weed species were used as training examples for a computer algorithm which is then used to discriminate all of the tree crowns detected in the imagery. The supervised learning approach allowed for a human user-in-the-loop to assign examples of different weed types in the collected flight data, resulting in better weed discrimination in varying weather and lighting conditions than a fixed set of weed detection/classification rules.

3.6.4 Final Classified Map Delivery

Once detections and classifications of weeds were made in the image data, these objects were transferred into the geo-referenced map, and displayed to the user on screen in different formats including a photo-mosaic of the area and a 3D surface model of the terrain.

4 Results and discussion

4.1 Processing Pipeline Timing

The time taken at each stage in the processing pipeline varied depending on the size of the area for which data was collected. Table 3 details approximate processing times for each of the steps required for building weed maps over a 1000x500m area. IMU and GPS data collected for a 1 hour flight amounted to approximately 500Mb of data which could easily be transferred to laptop via USB flash drive, whereas image data for a single flight typically amounted to over 30Gb, requiring portable hard drives to transfer data. The total time required for building a geo-referenced map was approximately 4 hours and was dominated by the time taken to process and match image features. The number of features extracted from each frame could be configured, with a reduction in features per frame resulting in a reduction in computation time and equivalently the accuracy of the final geo-reference map. The total time taken for tree crown detection and classification was in the order of 10 hours and was dominated by two factors. The first was the tree crown detection, which was a prototype version of software, which in principle should be able to run faster (in the order of 1 hour) with improved implementation. The second factor is the human-in-the-loop processing used to provide training examples of weeds to the classification algorithm. Typically we found this process took on the order of 1-3 hours for a detailed analysis on the ground. Most of this time was due to the operation of our user-interface which was not optimised; it is believed that future improvement of this interface could significantly reduce the time and effort required by the human-user.

Task	Time Taken	Notes
UAV Data Collection	60 minutes	
IMU/GPS Data Distribution	5-10 minutes	
Vision Data Distribution	2-3 hours	Requires ~30-60 minutes per computer, depends on other logistical factors
EKF Initial Pose Estimates	5 minutes	
Vision SIFT Feature Extraction	2-3 hours	
SIFT Feature Matching	~60 minutes	Depends on amount of image overlap
Map/Pose Optimisation	5-10 minutes	
Vision Tree Crown Detection	~8 hours	Current version is a prototype and future versions are expected to be significantly faster
Tree Crown Classifier Training	1-3 hours	Depends on human-used input, depends on number of training examples provided etc.
Tree Crown Classification	10 minutes	
Photo-mosaic generation	10 minutes	
3D Model Generation	~1-3 hours	Configurable for different levels of detail

Table 3 – Approximate processing times for different stages in map building and weed classification process

The total time taken to deliver weed classification results was 48 hours (operators needed to rest to be available to assist in flight activities during the day).

4.2 3D Terrain Mapping and Reconstruction Results

Figure A14 illustrates an entire area mosaic map build from the data collected during flight 10 of the UAV during the 2009 flight trials. The full resolution of the imagery available gives a spatial resolution of approximate 3-4cm per pixel over a 4000 by 600 meter area (low resolution image shown in this report due to space constraints). There are some small gaps observed in the coverage of the imagery due to wind disturbances acting on the UAV. Figure A15 illustrates a zoomed-in section of the mosaic map where different sections of Mitchell grass, woody weeds and some native trees are visible. Figure A16 illustrates a segment of constructed map of Williams Outstation from Flight 27 during the 2010 flight trials. Figure A17 shows a zoomed-in version of the same map. The maps have been reconstructed using the IMU, GPS and vision feature matches together in an optimisation stage. Typically maps were found to have horizontal positioning errors of 1-2m and vertical positioning errors of 2-3m globally (using comparison to ground reference points from GPS) but had much lower relative errors, visible through the high level of internal consistency seen in the maps.

Additionally, to the 2D photomosaic, 3D surface models of the terrain are also constructed, allowing for the visualisation of height in the environment. Figure A18 shows several oblique views of the flight area from Flight 27. The rough 3D structure (as viewed from above) and the heights of different trees are observable from the data. At this stage, the 3D height information is only used for visualisation purposes, however future work will explore the use of this information to assist in the segmentation and classification of different trees (see future work section below).

4.3 Tree Crown Detection Results

A tree crown detection algorithm developed at the ACFR was employed as a pre-processing step before classification of weeds in the map. This significantly reduced the complexity and time taken for classification by constraining the problem to that of classifying different trees, rather than having to classify background objects in addition to trees. Figure A19 illustrates example results on three different images of the first stage of tree crown detection where the image is segmented into regions of background, vegetation and shadow. The first stage worked reasonably well, however manages to miss-classify certain sections of the images (i.e. a large dirt patch in the first image has been miss-classified as vegetation due to its dark colour properties). The tree crown detection algorithm therefore employs a second stage which uses a tree and shadow image template to detect instances of tree and shadow combinations, with the correct shadow orientation inferred from an internal model of the time of day and sun angle. The result of the combined algorithm is shown in Figure A20 where objects are detected based on the correlation of the template model with the original segmented image. Various template sizes are used for matching to detect trees of varying sizes (3 different sized tree templates are used currently).

4.4 Weed Classification Results

Figure A21 illustrates a section of the flight map in flight 11 taken at the Carrum Farm site in 2009 with classified vegetation data overlaid in the map. In this case the classifier has labelled each section of the map as one of four classes: Prickly Acacia (shown in blue), Parkinsonia (shown in red), Eucalyptus trees (shown in yellow) and Mitchell grass/other (not shown on map). The classification algorithm finds large populations of Prickly Acacia surrounding the dried-out riverbed along with a small number of sparse Parkinsonia bushes and Eucalyptus trees. Figure A22 illustrates a zoomed-in view of the classified map with samples of the ground survey data (shown by the circles) also overlaid in the image. The ellipses surrounding the ground survey points indicate the 5m confidence in the measurements (based on the accuracy of the handheld GPS receiver).

Figure A23 illustrates the location and classification of tree crowns detected in the flight area segment of Flight 27, taken during the 2010 flight trials. Figures A24 and A25 illustrate zoomed-in views of different areas in the map including classifications. In this flight area, all tree crowns were separated into four different classes; two weeds of national significance (Parkinsonia and Mesquite), a smaller, less troublesome weed (Mimosa) and a class representing a variety of Eucalyptus and Coolabah native tree species found in the area (Native). The classes were selected based on tree types discovered during a ground-survey of the area. Several different species of native trees were found, and these were grouped into one general class based on their appearance.

Table 4 shows classification results provided by cross-validation based on a selection of training examples provided by ground-surveys of the flight area. Cross validation was performed by choosing one of the tree crowns, removing this crown along with another 10% of the crowns (randomly selected) and training the classifier using the remaining tree crowns with labels provided by the ground survey. The trained classifier was then used to classify the originally removed tree crown, and the result recorded as either a correct or incorrect classification. This process was repeated by cycling through each of the tree crowns provided in the training data. The performance of the classifier was then measured via three metrics; accuracy, precision and recall. The overall accuracy was calculated as the ratio of correct classifications to total tree crowns. For each class, we evaluated the precision and recall in classification. Precision is defined as the percentage of correct classifications of a class compared to the total number of objects examined in that class. Recall is defined as the ratio of correct classifications to the sum of the correct classifications and incorrect classifications of the given class as other classes (i.e.

missed classifications of the original class). Finally, a confusion matrix was drawn up which represents the number of correct/incorrect classifications of each class. The diagonal values of the confusion matrix represent the number of times a given class was correctly classified whereas the off-diagonal terms represent the number of times a certain class was miss-classified as an alternate class (and which class it was miss-classified as).

		Actual			
		Parkinsonia	Native	Mimosa	Mesquite
Classified	Parkinsonia	76	8	2	57
	Native	11	268	3	31
	Mimosa	2	3	3	7
	Mesquite	69	39	12	573

Confusion Matrix

	Parkinsonia	Native	Mimosa	Mesquite
Precision	48.10%	84.28%	15.00%	85.78%
Recall	53.15%	85.62%	20.00%	82.68%
Overall Accuracy	79.04%			

Classification Results via Cross-Validation

Table 4 – Summary of tree classification results via cross-validation for Flight 27: Shown is the confusion matrix and precision/recall values for each of the four classes: Parkinsonia, Native tree, Mimosa and Mesquite.

The classifier displays high scores in both precision and recall of the native species and mesquite (~85%). These two classes were reasonably distinguishable based on the colour and texture features used in classification. The precision and recall values displayed for the Parkinsonia and Mimosa were somewhat lower. In the case of mimosa bush, the classification result was poor due to the low number of training examples provided in the data. This was due to the relatively low amount of mimosa bush present in the area and the small size of the bush, meaning that there was a higher chance it would be missed during the tree crown detection phase, resulting in a low number of training examples. Secondly, the colour and texture features of Mimosa were very similar to Mesquite, making it difficult to distinguish (as seen by the confusion matrix, Mimosa was most often miss-classified as Mesquite). The classification results for Parkinsonia were also relatively poor compared to Mesquite and the native tree; examination of the confusion matrix shows that when miss-classified, Parkinsonia was more often mistaken for Mesquite than other classes. This is likely due to the relatively high distinguishability of the native trees from weeds based on their colour and texture properties in the images.

Overall, the results demonstrate a very good ability for distinguishing weeds such as Mesquite and Parkinsonia from native species (such as Eucalyptus and Coolabah) but a reduced ability to distinguish weeds from one another. This is observed also in the classified map results in Figures A23 to A25; the results often display a mix of different weeds species identified in a single tree crown.

5 Success in achieving objectives

The project and the flight trials at Julia Creek during 2009 and 2010 were considered a success based on the following achievements:

(1) Demonstration of repeated data collection using the FUAV system
(2) Demonstration of fast mapping and classification of weeds (24-48 hours) using simple computing resources (laptop computers)
(3) Flight demonstration of the RUAV system

The ability to perform data processing in the short time span using simple computing resources, demonstrated the feasibility of a data processing system that would be cost effective (could run on laptops or standard home desktop computers) without the need for extensive or specialised computer resources (i.e. computers servers). The quick turn-around time also supports the ability to supply information to land managers for immediate use and to use for autonomous spraying via the HUAV system. Classification results demonstrate the potential for semi-autonomous classification of different weed species, however some work is needed to improve on the accuracy of this result. Due to relatively strong winds we were not able to demonstrate the complete functionality of the RUAV system dispensing the herbicide as a complete system. Instead the RUAV was demonstrated separately from the auto-dispensing system. The results from those trials demonstrated that a little more work is required to complete the system for greater robustness. This does not need to be undertaken again at the farms but could be completed at the University property.

6 Impact on meat and livestock industry – Now and in five years time

The successful outcomes from this project have demonstrated the potential for the use of UAVs and weed detection/discrimination algorithms in several areas of agriculture, farming and the meat and livestock industry.

6.1 Reduced Cost using UAVs for On-farm Management

The UAV system has the potential to be cheaper and safer to run than manned aircraft or ground-based weed management activities. The demonstration of both RUAVs and FUAVs within the farmland environment in this project has established the potential benefits for UAVs in other on-farm management activities such as monitoring crop growth, water usage, wear and tear on gates and fences and monitoring of livestock over large areas. Combining these elements into a single system has the potential to produce cost savings in other areas of farm management.

6.2 Weeds Discrimination and Classification Algorithms

The classification and mapping algorithms developed in this project are able to distinguish between weeds and native species rapidly and with minimal time and labour requirements, covering extensive areas with large amounts of data. These algorithms also have the potential for use in other remote sensing contexts outside of UAVs such as with manned aircraft or possibly satellite imagery.

6.3 Automated Precision Herbicide Dispensing System

The development of the RUAV and auto-herbicide dispensing system has demonstrated the ability for automated and precision delivery and could potentially be used separate of the weed discrimination algorithms.

7 Conclusions and recommendations

This section of the report summarises the findings of the project and details directions for future research in the area.

7.1 Role of UAVs in Weeds Monitoring

This project has demonstrated the feasibility of a UAV system for mapping and classifying various species of woody weed such as Prickly Acacia, Mesquite and Parkinsonia over small-scale areas using a low-flying UAV system carrying a low-cost sensor payload consisting primarily of a 3-band colour vision camera. Previous and current approaches to woody weed mapping use remote sensing data from satellite imagery [2] or manned aerial surveys [1], providing imagery resolutions of at most 0.5-1.5m/pixel; the imagery resolutions achieved in this project were on the order of 3cm/pixel making it the highest resolution survey of its kind. The use of high resolution mapping has allowed for detection and classification of woody weeds at a very fine scale (individual shrubs), not yet achieved using existing remote sensing technology. Unlike satellite and manned-aerial surveying, the use of a low-flying UAV limits the coverage area achievable for surveying; the fixed-wing platform used in surveying was able to map a 240 hectare area in a single 1-hour flight.

The use of UAVs provides a number of advantages over manned aerial surveys. Data collection flights were largely automated, with a remote-human pilot required only during takeoff and landing of the UAV. The small size of the UAV meant that flights could be performed in remote areas, not requiring any special infrastructure for take-off and landing, and that fuel and maintenance costs were potentially much lower than manned aircraft. Future advances in UAV autonomy will increase the value in these systems for monitoring and mapping through avenues such as multi-UAV mapping systems and potentially smart, adaptive data collection strategies (see discussion below). These advances are one avenue for increasing the area coverage ability of these systems, and thus their value as a mapping tool when compared the manned aerial surveys and satellite imagery.

7.2 Future Work in UAV Ecology Mapping

The vision-based mapping and classification algorithms developed during this project have demonstrated the ability to build high resolution 3D maps using low-cost sensor equipment such as monocular vision. The ability to infer 3D structure in the environment opens up the possibility of a range of applications such as tree height estimation and has the potential to assist in decisions of tree classification, growth and age. Several advances were made during the project in the processing time required to construct maps. One avenue for future research is to work towards real-time map construction on-board the UAV platform during data collection. Real-time mapping (and potentially real-time weed classification) would enable online intelligent and adaptive data collection strategies (ideally suited for a UAV) where the UAV could make intelligent decisions about flight altitude and area coverage, maximising the information contained in the collected data, and allowing for more effective data collection.

7.3 Future Work in Weed Classification

The classification results presented in this project demonstrate a very good ability for distinguishing weeds from the background landscape and native trees and vegetation, but a limited ability to distinguish between different weed species (i.e. Mesquite and Mimosa). One avenue of future research is through the use of multi-spectral sensing from a UAV to improve distinguishability in the collected data. Multi-spectral sensing is becoming an affordable option (based on the cost and weight of these sensors) for small UAVs; near-infrared band imagery is a well-known source for distinguishing vegetation in aerial and satellite remote sensing and could potentially add to the accuracy of high resolution survey classification. Another avenue for improved classification (inspired by the discussion in [2]) is through the use of seasonal change detection, through multiple UAV flights, where changes due to flowering and leaf growth/loss could aid in distinguishing different weeds. So far only colour and texture properties in the vision data have been used as features to assist in classification; additional properties such as tree size, shape and height (estimated using 3D map data) could also help to assist in tree classification using existing vision data; this is an area of future research.

7.4 Future Work in Hovering UAV System/UAV-based Herbicide Delivery

Although the RUAV herbicide delivery system was not demonstrated as a whole, each component of the system, including the RUAV, autonomous flight operation and the operation of the granule delivery system were tested successfully, thus demonstrating the feasibility of the approach. Based on the weed location accuracy achieved from the fixed-wing UAV maps, a recommended line of future work would be to include a vision sensor on the hovering UAV to act as an extra source of weed tracking information for final stage discrimination of weeds during spraying. This would allow for sub-meter positioning accuracy for herbicide delivery through real-time final stage tracking of weeds from the hovering platform, potentially resulting in improved effectiveness of the herbicide, and increased robustness.

7.5 Analysis of Cost and Requirements for Up-Scaled System Operation

Based on the results of the demonstrated multi-UAV system for weed mapping and herbicide delivery, an analysis was performed into the feasibility and associated costs for applying the same type of UAV system to a larger 10-by-10km area (10,000 hectares). Based on the existing FUAV and sensor payload configuration, the coverage area possible for a single flight is primarily a function of the flight altitude and is shown in Table 4. Scaling-up the flight patterns to cover a 10-by-10km area would thus require several flights, depending on the altitude and is also shown in Table 4.

Existing Flight Patterns	
Flight Altitude	**Area Covered in 1 hour**
100m	4000-by-600m
500m	6000-by-1500m

Requirements for mapping 10-by-10km area	
Flight Altitude	**Number of Separate 1-hour Flights**
100m	40
500m	8

Table 4 – Existing flight area coverage capabilities and number of flights required for 10-by-10km area coverage using the fixed-wing UAV system

The existing system relies of a flight altitude of 100m, thus resulting on a ground imagery resolution of 3cm/pixel for the existing 1 megapixel camera used in the sensor payload. To make the system more feasible, it is recommended that this be replaced with a 12 megapixel camera (which is commercially available), thus allowing the system to fly at an altitude of 500m while achieving approximately the same ground resolution as the existing system does at 100m, in turn allowing for a larger area coverage in a lower number of flights. The resulting system would thus have the following potential mapping characteristics shown in Table 5.

UAV Flight Parameters	
Flight Speed	25 m/s
Flight Altitude	500m
Sensor Resolution	12 megapixels
Ground Resolution	~3-4cm/pixel

Requirements for Mapping a 10-by-10km area	
Number of Flights	8
Time Per Flight	1 hour
Number of Flight Days	2-3 days
Data Processing Time per Flight Hour	4-8 hours per flight
Total Data Processing Time	32-64 hours

Table 5 – Proposed fixed-wing UAV mapping system for 10-by-10km operation based on the analysis of results from the existing UAV system

The times associated with flying and data processing are approximate and are based on the processing times achieved in this project with existing flight experiments. The 2-3 days of flying time accounts for the fact that optimal flying time is restricted to a few hours per day (based on lighting conditions) and accounts for time taken between flights due to operational requirements, such as downloading data and refuelling the UAV.

Based on the operational system described above, an approximation of the cost associated with the operation of this system is provided in Table 6.

Up-front System Costs	
Fixed-wing UAV Platform	~$10,000
UAV Sensor Payload	~$20,000
UAV Ground Station Equipment	~$5,000
Data Processing Facilities	~$5,000

Ongoing Costs/ Cost per 2-3 day flight operations	
UAV Fuel Costs	~$80 (5L/flight, $2/L)
UAV Maintenance Costs	~$240 (4 hours maintenance over 3 days)
Number of Required Personnel	2-3

Table 6 – Analysis of costs associated with 2-3 day flight operations required to cover a 10-by-10km area for weed mapping

In addition to the upfront costs shown in Table 6, two engineers would be required to build the UAV platform and install the sensor payload (one aeronautical engineer and one electronic/computer engineer) over approximately a one-year period, and thus the salaries of these staff would need to be accounted for. Based on the operational experience during flight trials at Julia Creek, approximately 2-3 personnel would be required to run the system, including a remote-control pilot (for takeoff and landing and for safety reasons) and a payload specialist for operating the ground station and managing data coming off the UAV system. The salaries of these operating staff would also have to be factored into operation costs.

8 Bibliography

[1] R.D. Klinken, D. Shepard, R. Parr, T.P. Robinson and L. Anderson (2007). "Mapping Mesquite (Prosopis) Distribution and Density Using Visual Aerial Surveys", Rangelands and Ecology Management, 60(4):408-416.

[2] J. Muir (2010). "Using Remote Sensing to map the Invasive Weed Prickly Acacia (Acacia nilotica ssp. indica) in the Mitchell Grass Downs Bioregion". Australian Remote Sensing and Photogrammetry Conference, Alice Springs, 2010.

9 Appendices

9.1 Flight Platform and Sensor Payload System

Figure A1 – J3 Cub UAV System in flight over Julia Creek, Queensland during the July 2009 flight trials. The platform and sensor payload have remained the same for the 2010 flight trials.

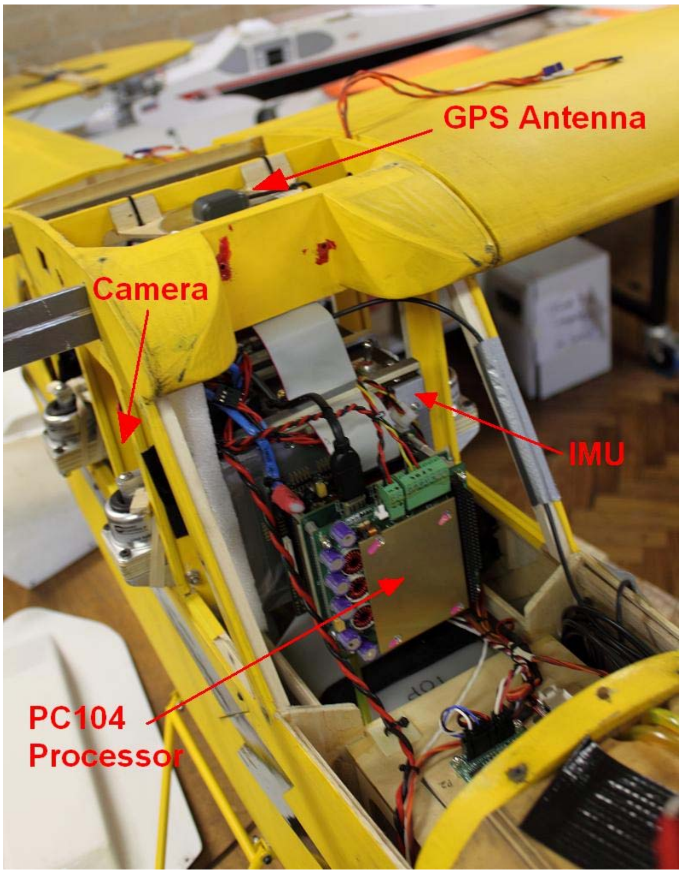

Figure A2 – Internal view of the J3 Cub UAV System showing payload box incorporating the IMU, GPS receiver and antenna, vision-stack processor and camera system.

Vision Camera	Hitachi HV-F31	**IMU**	Honeywell HG1900
Sampling Rate	3.75Hz	Sample Rate	600Hz, pre-processed to 100Hz
Field of View	28^o x 22^o	Accel. Noise (1σ)	$0.05 m/s^2$
Resolution	1024 x 768 pix	Gyro Noise (1σ)	$0.05^o/s$
Angular Resolution	0.0285^o	Accel. Bias(1σ)	$0.05 m/s^2$
Ground Resolution	3.7cm/pix @ 100m, 18.6cm/pix @ 500m	Gyro Bias (1σ)	$0.05^o/s$
Ground Footprint	38x30m @ 100m, 190x150m @ 500m		

GPS Receiver	Novatel OEM5, differentially corrected
Sample Rate	5Hz
Position Error (1σ)	1m
Velocity Error (1σ)	10cm/s

Table A1 - Sensor Payload Specifications: The sensor payload consists of an IMU, GPS receiver and downwards-mounted colour monocular camera.

Figure A3 – Hovering UAV System in flight over Julia Creek, Queensland during the July/August 2010 flight trials.

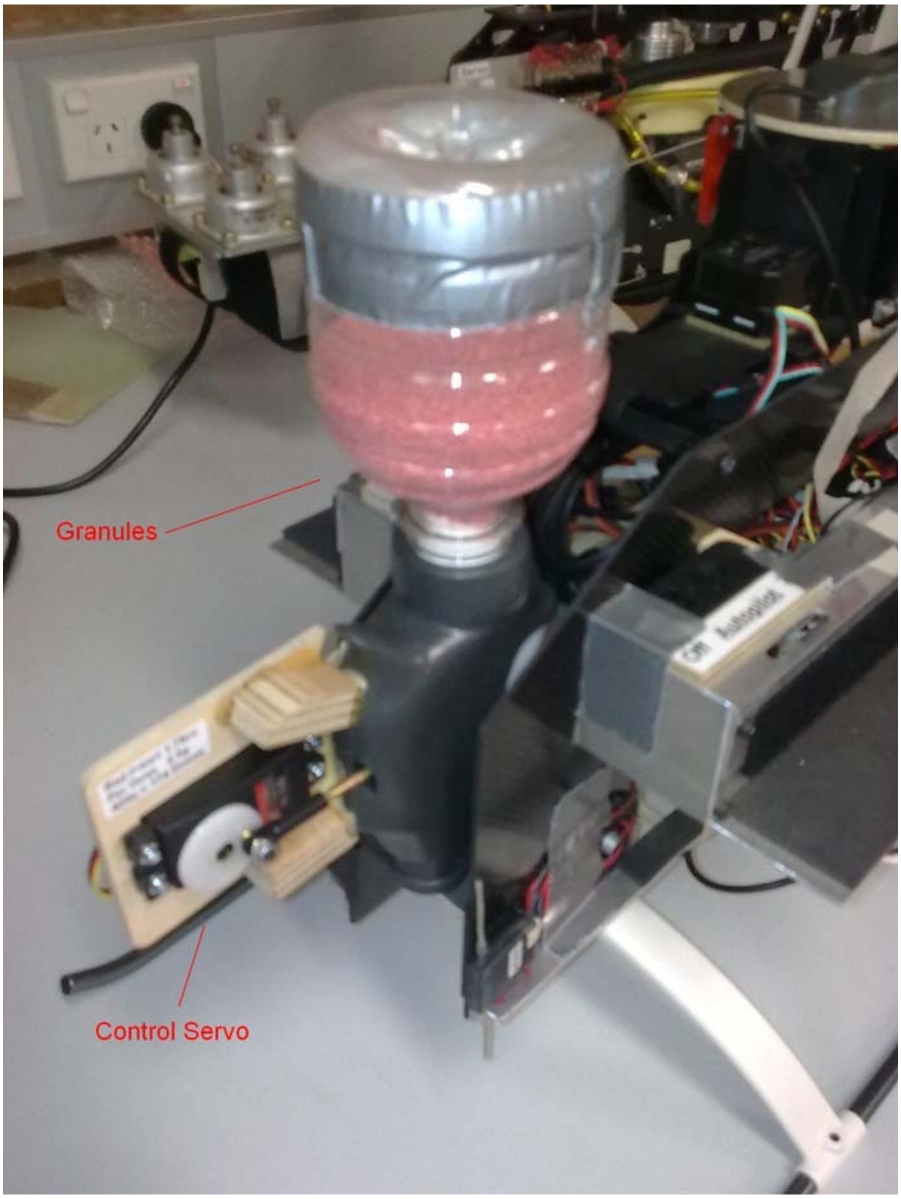

Figure A4 – Herbicide Granule Dispenser on the hovering UAV with remote control activated servo for controlled granule delivery

9.2 Mapping and Classification Algorithms

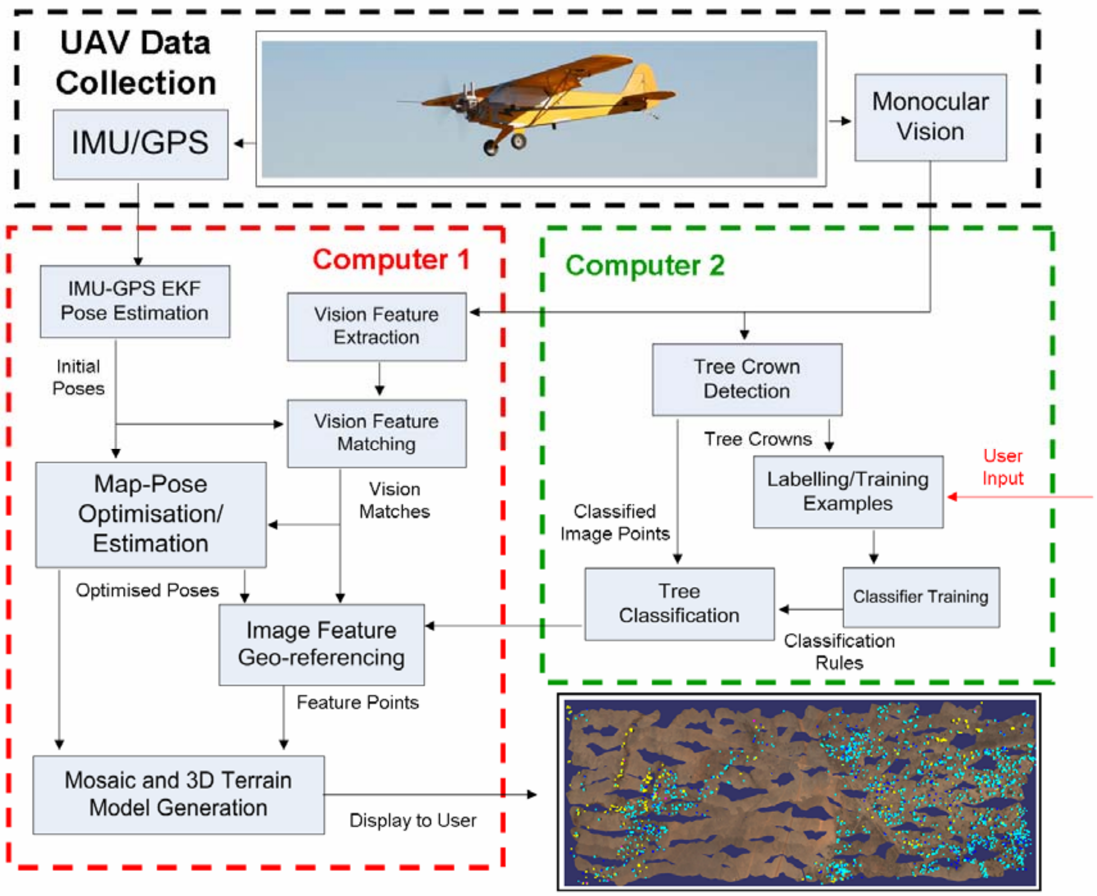

Figure A5 – Processing Pipeline for Weed Mapping and Classification Algorithms: IMU, GPS and Vision data are copied from the UAV and distributed to two laptops computers. Computer 1 builds a geo-referenced map of the area using IMU/GPS and vision data. Computer 2 performs tree crown detection and classification using the images. Detected and classified weeds are then integrated into the map and displayed to the user.

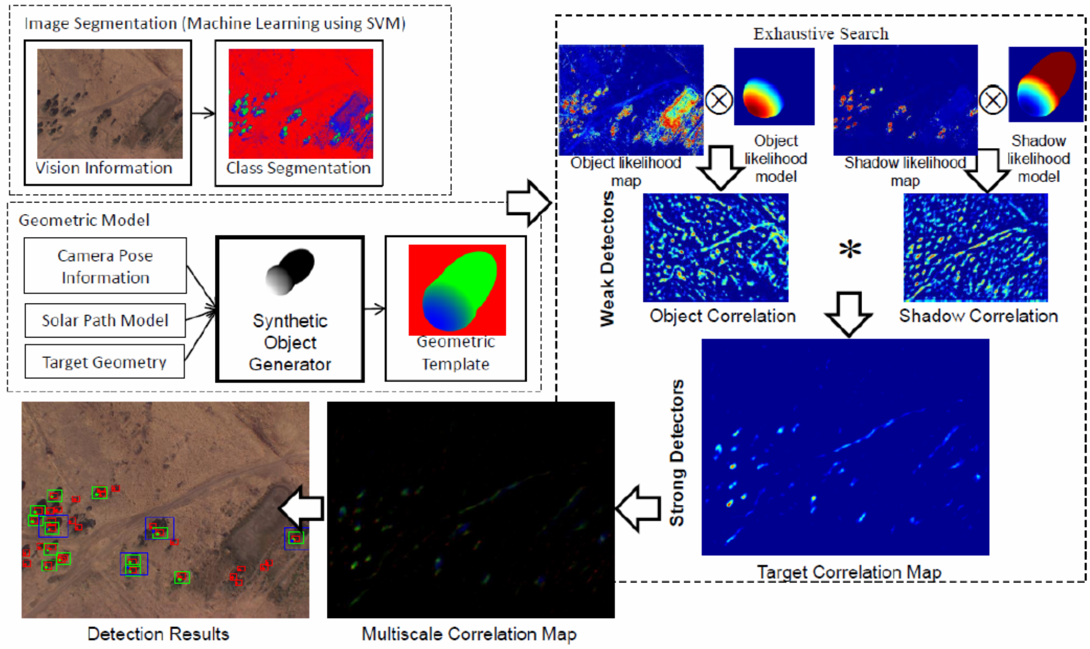

Figure A6 – Processing Pipeline for Tree Crown Detection Algorithm: Each image is segmented into three classes: background, tree and tree shadow. A second stage detects tree crowns using template matching of a tree/shadow object using information about the time of day and sun angle.

9.3 2009 Julia Creek Flight Trials

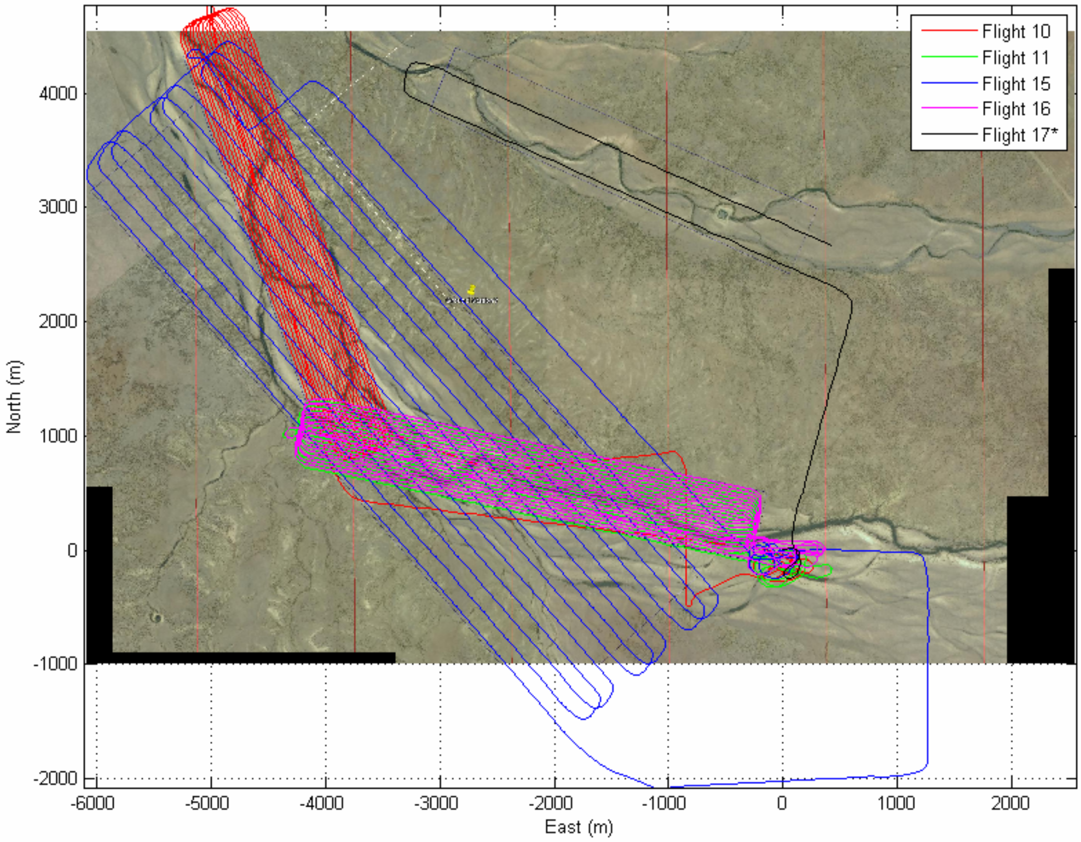

Figure A7 – Flight paths for Flights 10, 11, 15, 16 and 17 at the Carrum farm site. Shown underneath the flight paths is low resolution map imagery of the area available from Google Earth. (*Flight 17 is only partially displayed due to an error in the collected GPS data for this flight).

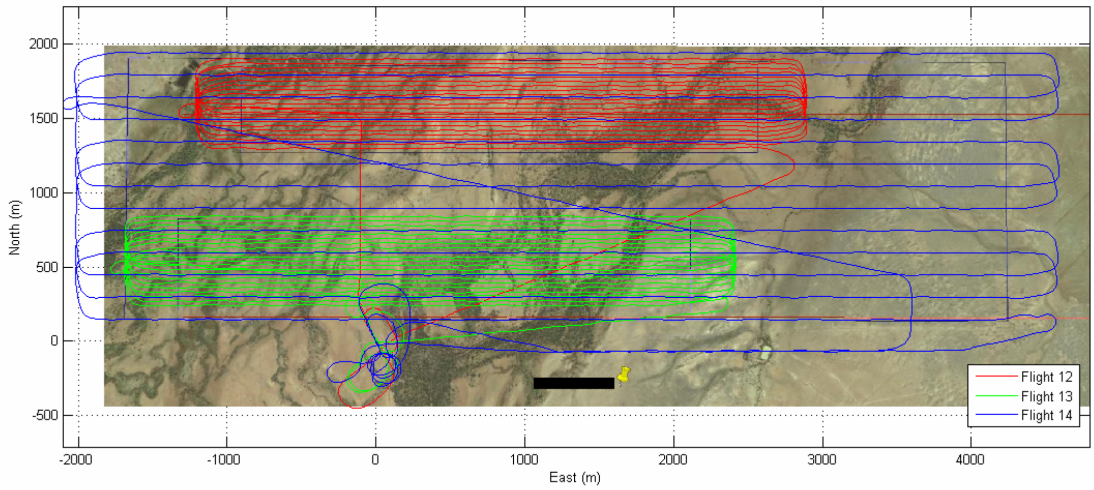

Figure A8 – Flight paths for Flights 12, 13 and 14 at the Williams outstation site. Shown underneath the flight paths is low resolution map imagery of the area available from Google Earth.

Figure A9 – Ground based photographs of different woody weeds located in the Julia Creek flight area: left, Prickly Acacia and right, Parkinsonia.

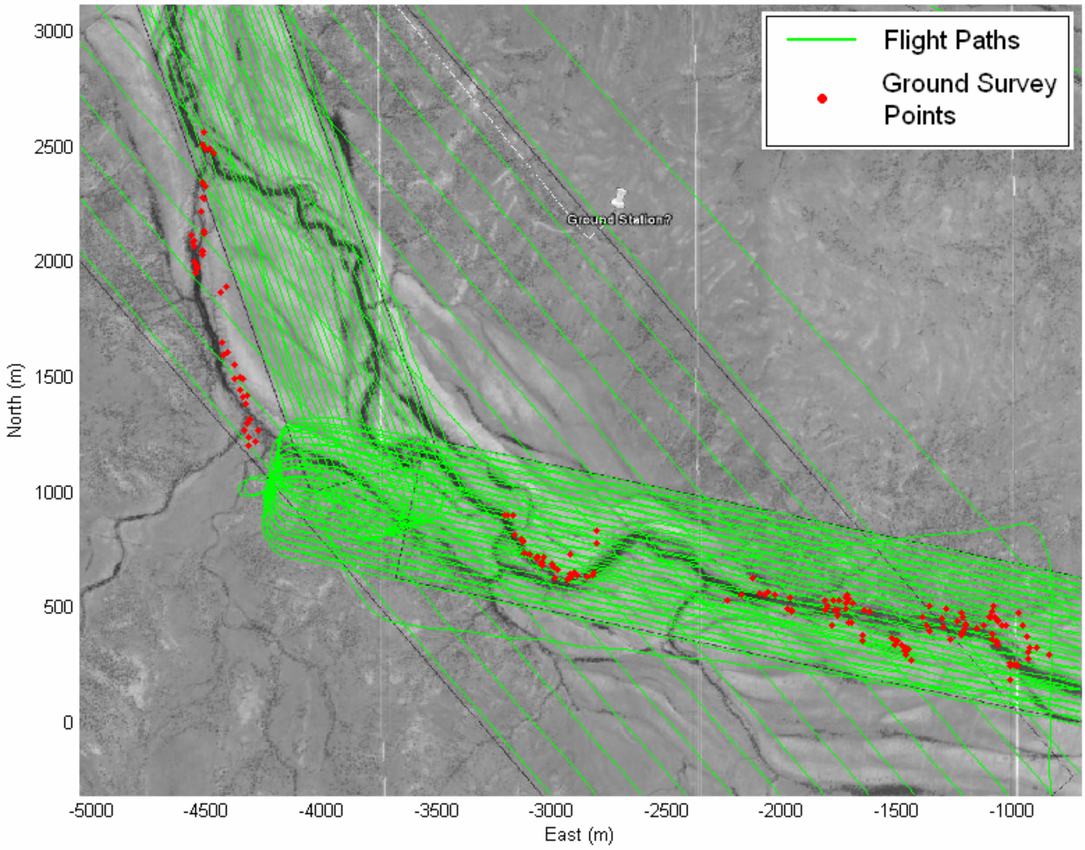

Figure A10 – Ground Survey Data Collected at Carrum Farm Site: The green lines indicate the paths of different flights in the Carrum area where the red points indicate the locations of surveyed woody weeds and other native vegetation.

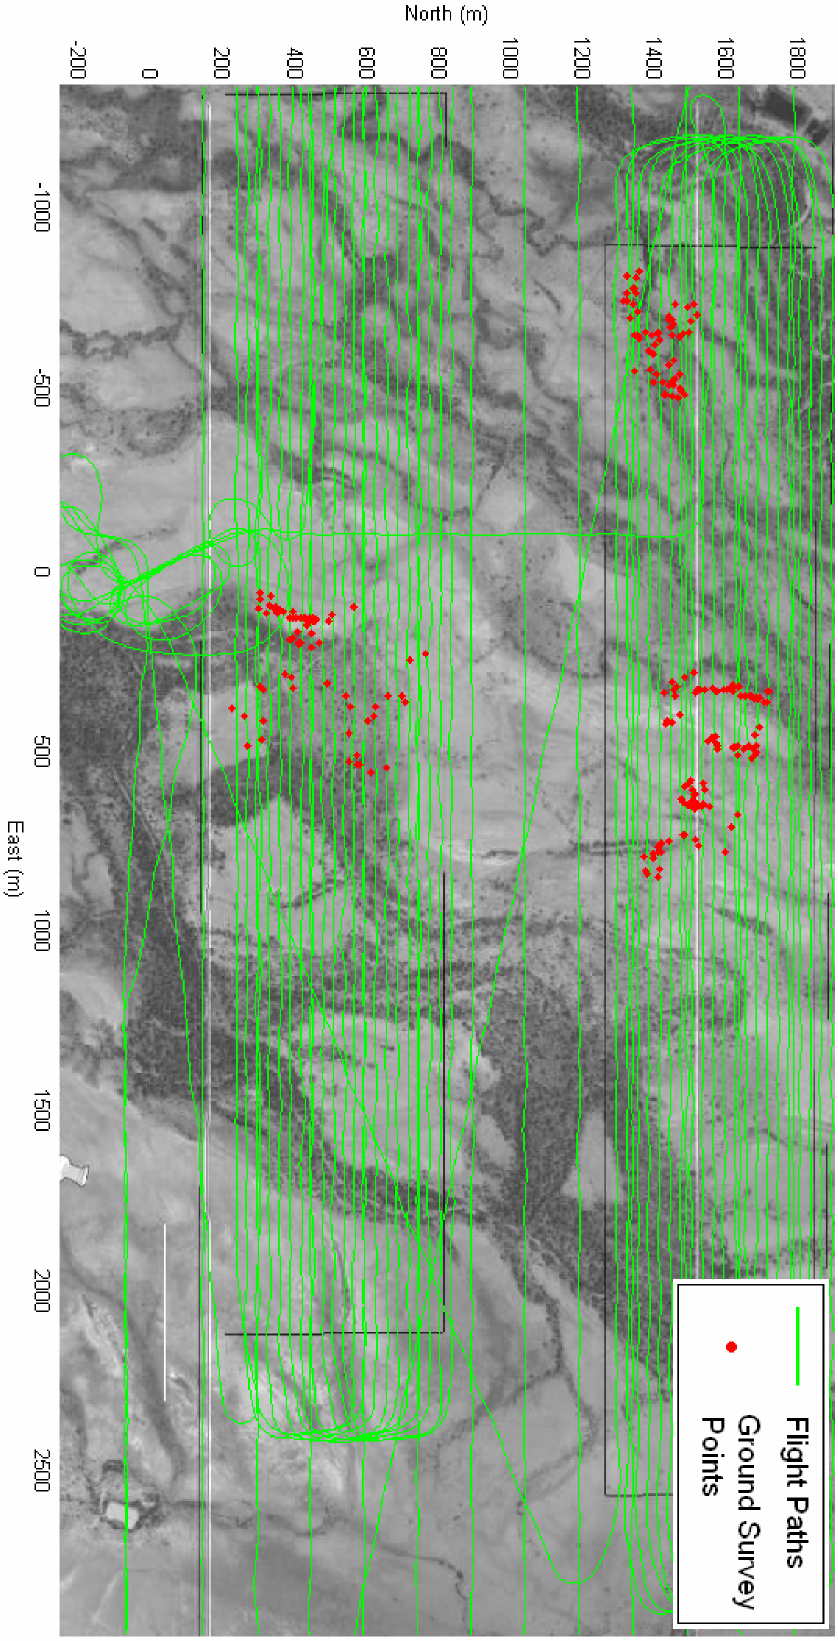

Figure A11 – Ground Survey Data Collected at Williams Outstation Site: The green lines indicate the paths of different flights in the Carrum area where the red points indicate the locations of surveyed woody weeds and other native vegetation.

9.4 2010 Julia Creek Flight Trials

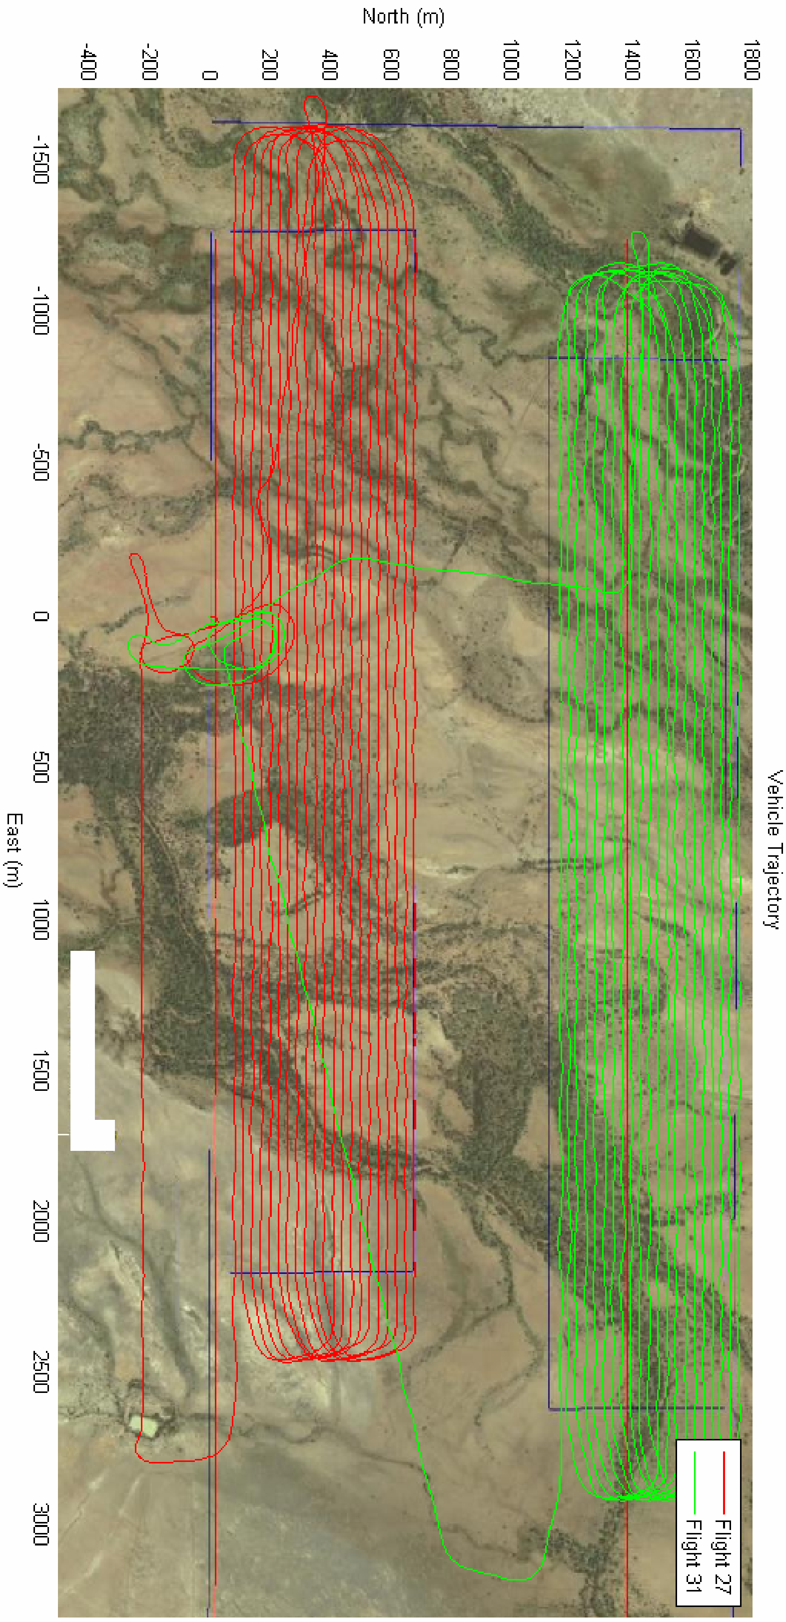

Figure A12 – Flight paths for Flights 27 and 31 (Flight 28 followed same flight path at Flight 31) at the Williams Outstation site. Shown underneath the flight paths is low resolution map imagery of the area available from Google Earth.

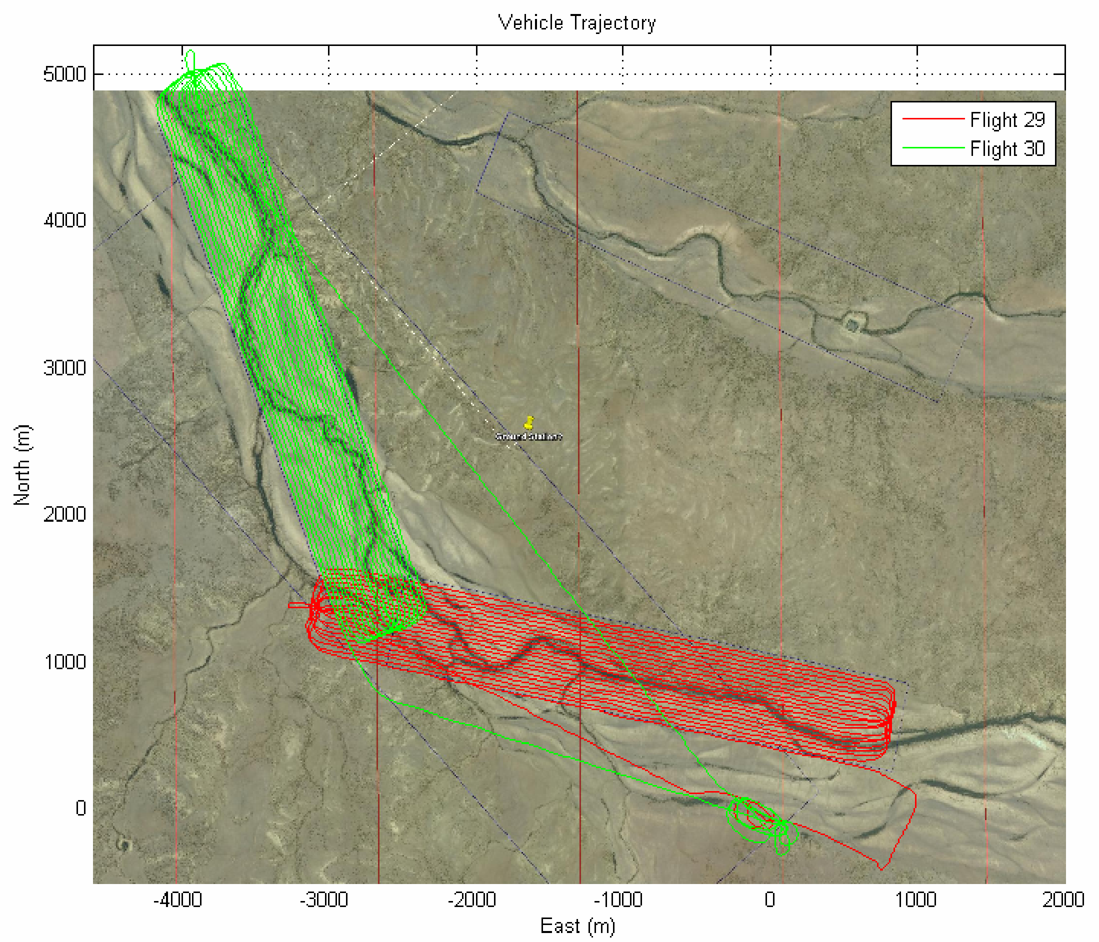

Figure A13 – Flight paths for Flights 29 and 30 at the Carrum Farm site. Shown underneath the flight paths is low resolution map imagery of the area available from Google Earth.

9.5 Mapping and Classification Results

Figure A14 - Final Environment Mosaic for Flight 10: Mosaic map of the entire flight area for flight 10 covering a distance of approximately 4000 by 600 meters at the Carrum Farm site.

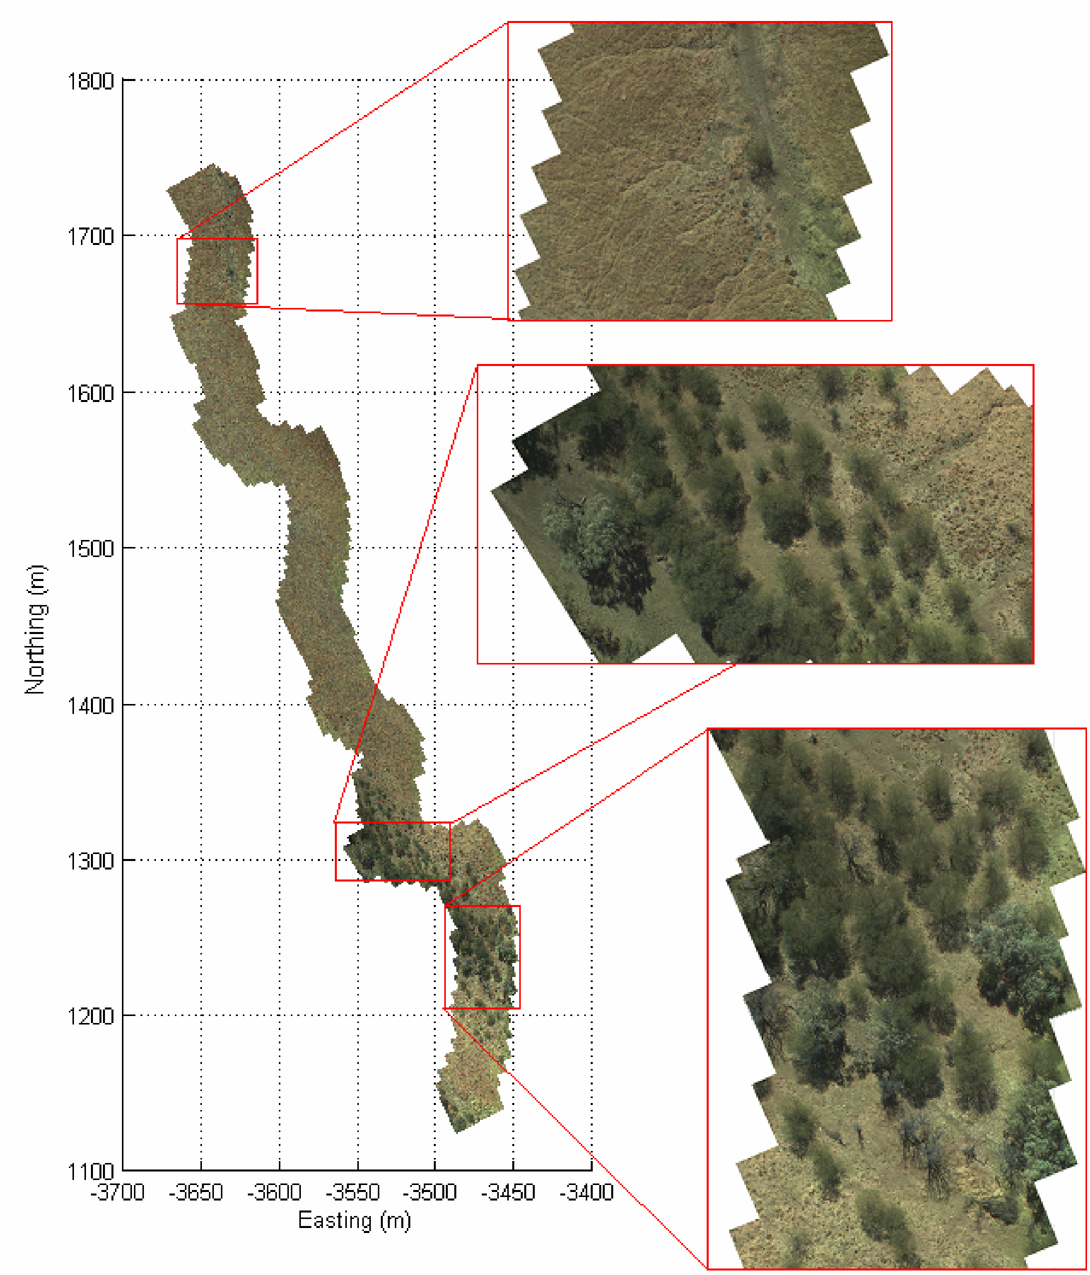

Figure A15 – Zoomed-in section of the mosaic map for flight 10 showing areas covered by grass, trees and different types of vegetation. The terrain mapping and mosaicing systems produce a consistent, geo-referenced map of the environment which can be used for classification and visualisation of the environment.

Figure A16 – Segment of constructed map of Williams Outstation from Flight 27 (gaps seen due to platform motion, future work will look at online data collection planning for improved area coverage)

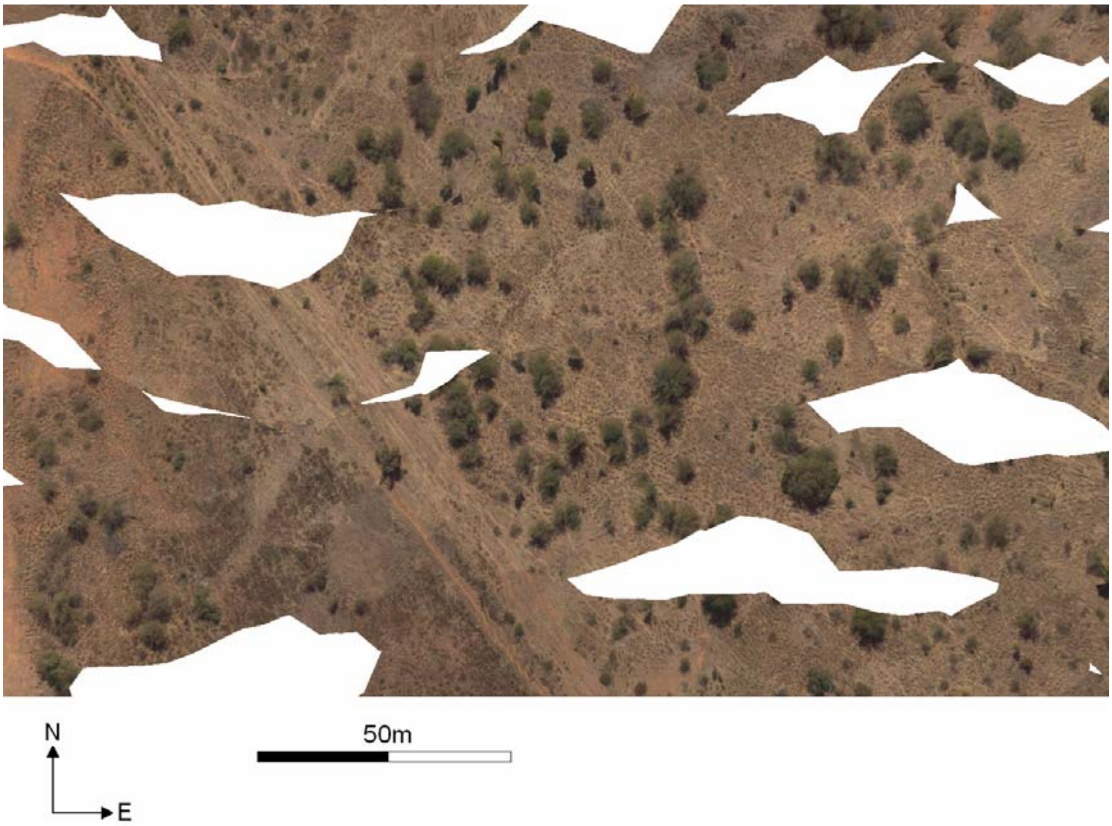

Figure A17 – Segment of constructed map of Williams Outstation from Flight 27 (zoomed-in)

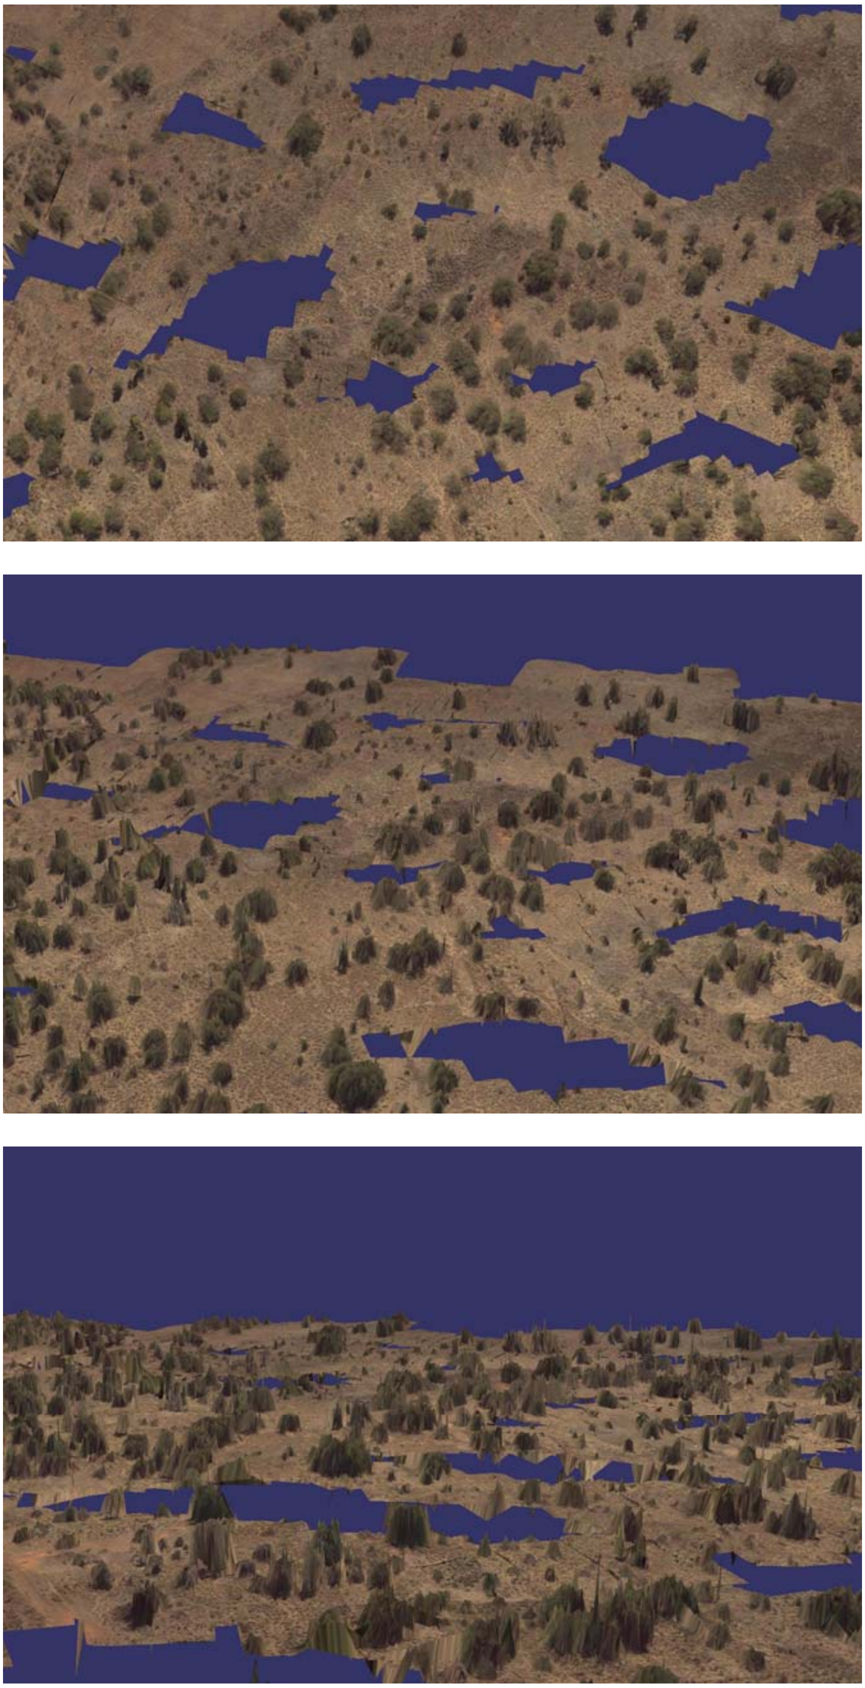

Figure A18 – Different oblique views of the 3D terrain map constructed from IMU, GPS and vision feature match data. The 3D map contains the coarse structure of vegetation, which could potentially be used as part of the classification algorithms (see future work section).

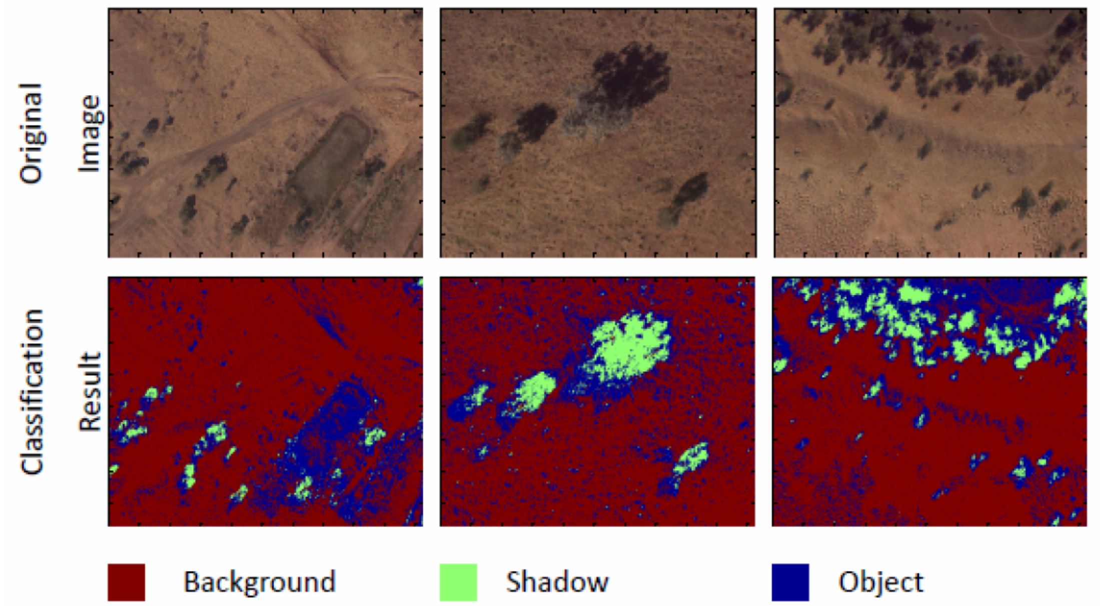

Figure A19 – Example results of first stage image segmentation for tree crown detection

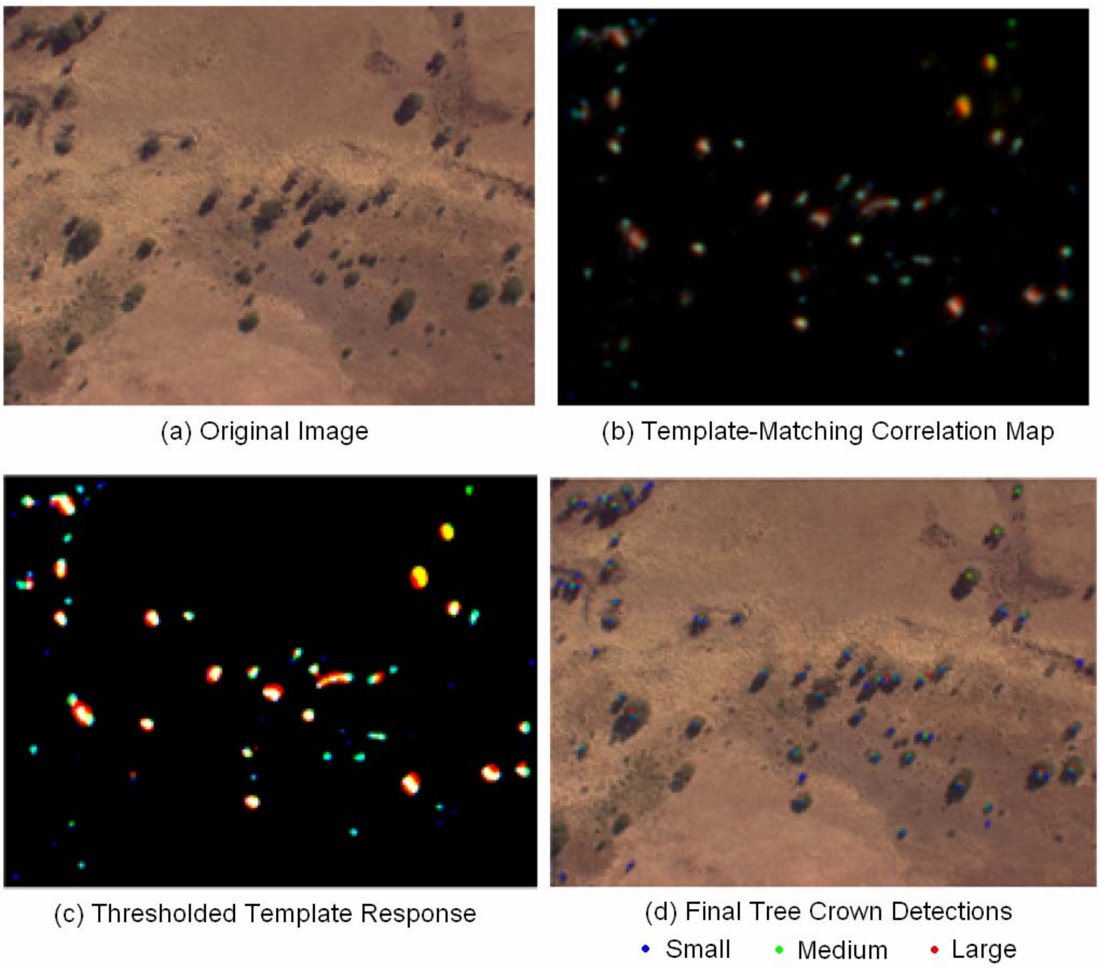

Figure A20 – Example results of second stage object/shadow template matching for tree crown detection

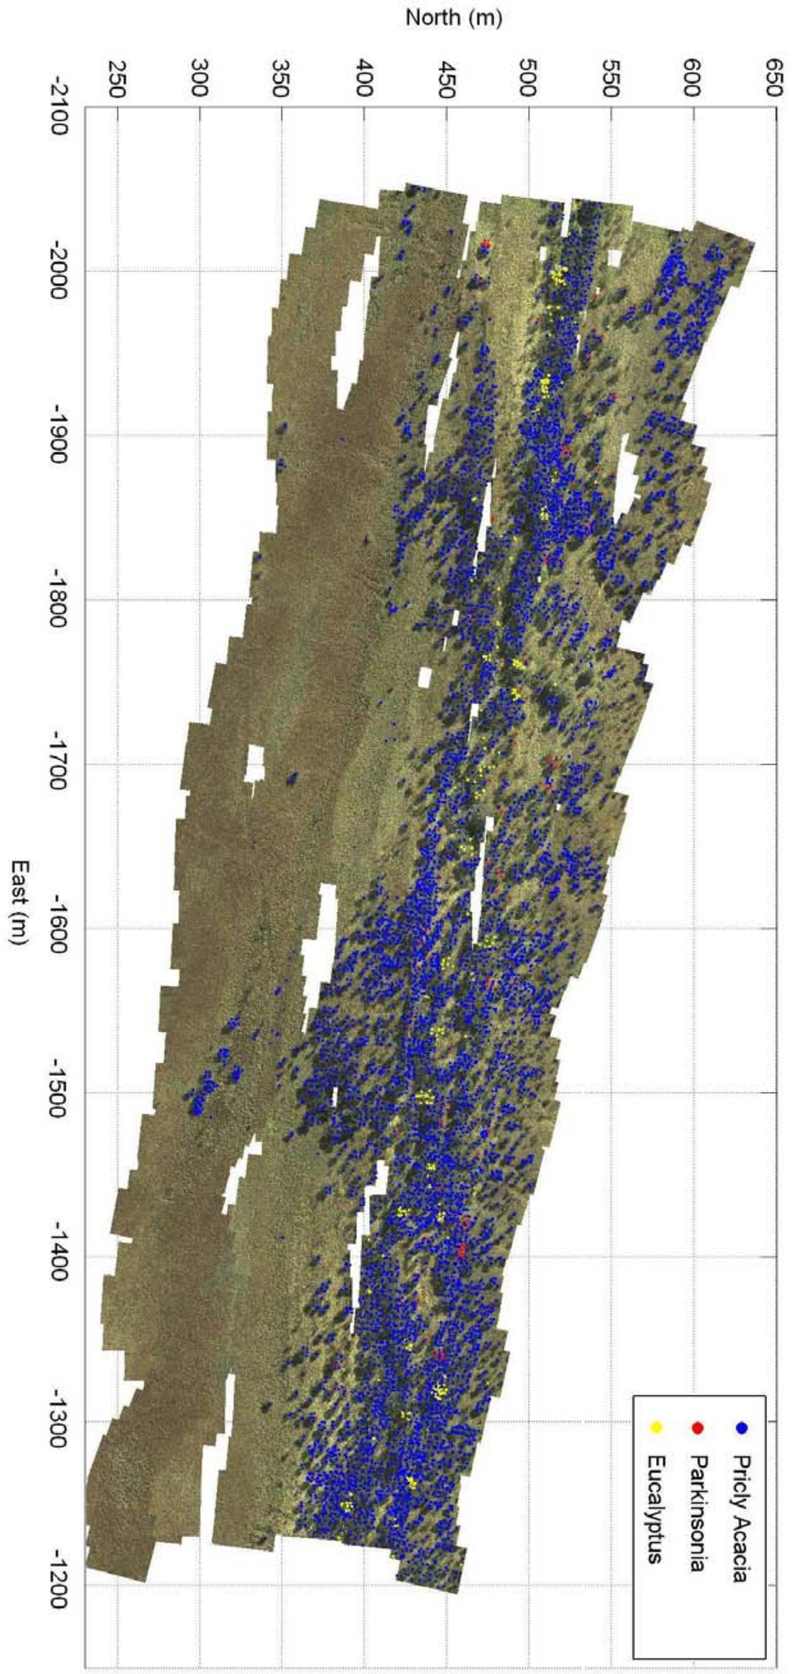

Figure A21 – Section of the raw photo-mosaic constructed for the flight area at the Carrum farm site using data collected during flight 11 with overlaid classification data for woody weeds Prickly Acacia and Parkinsonia and the native species of Eucalyptus.

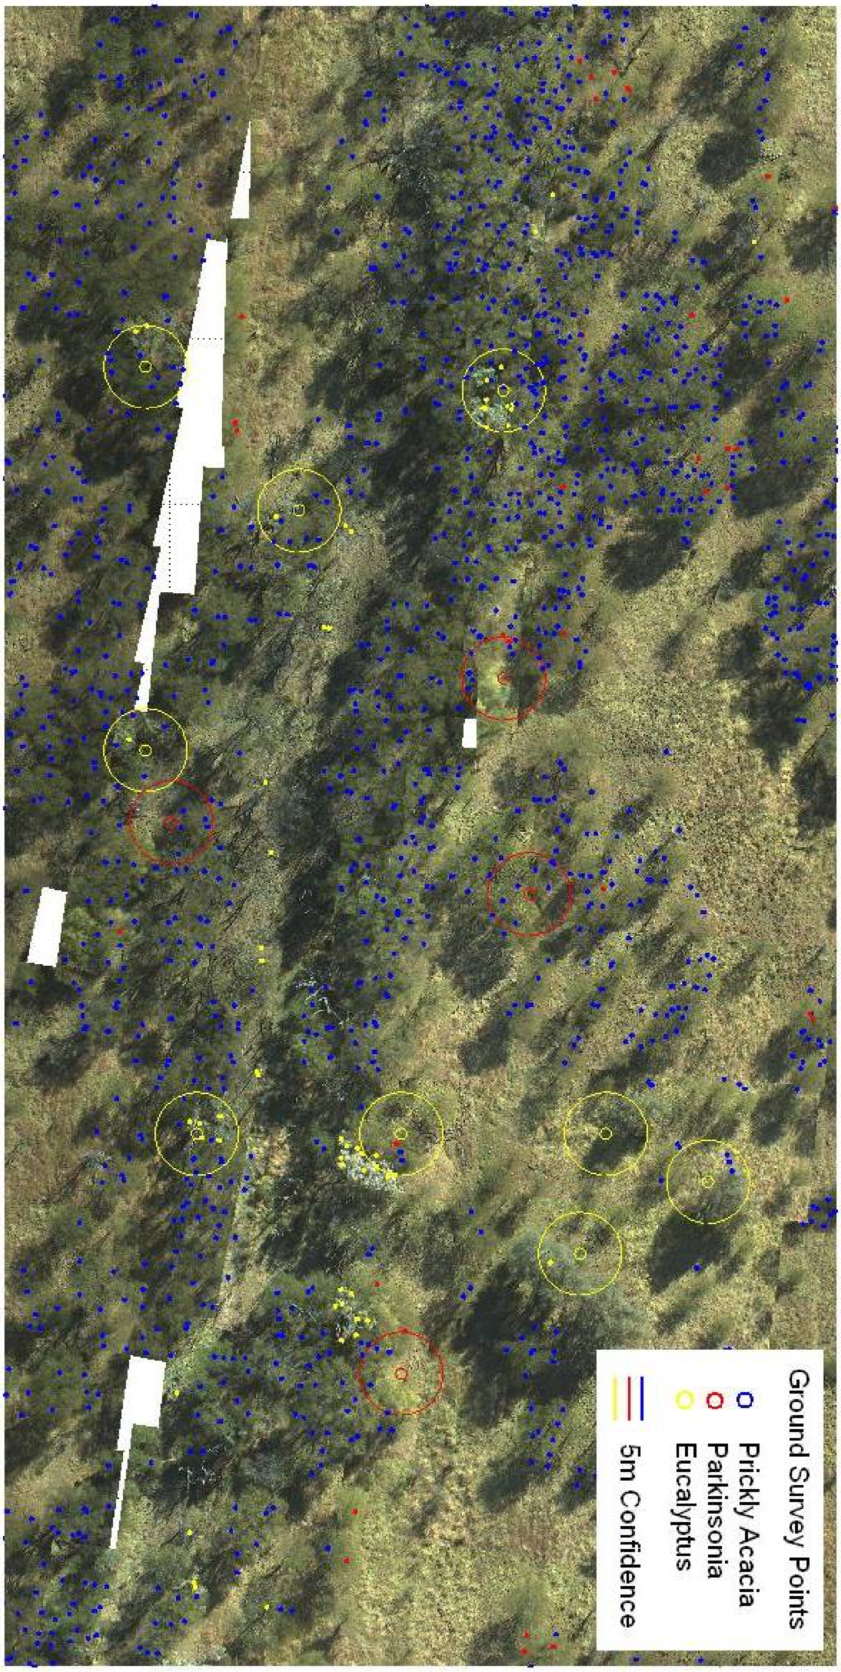

Figure A22 – Zoomed-in section of mosaic and classified vegetation from flight 11 data. Also illustrated are samples of the collected ground truth data which was used to help train the test the weed classification algorithms.

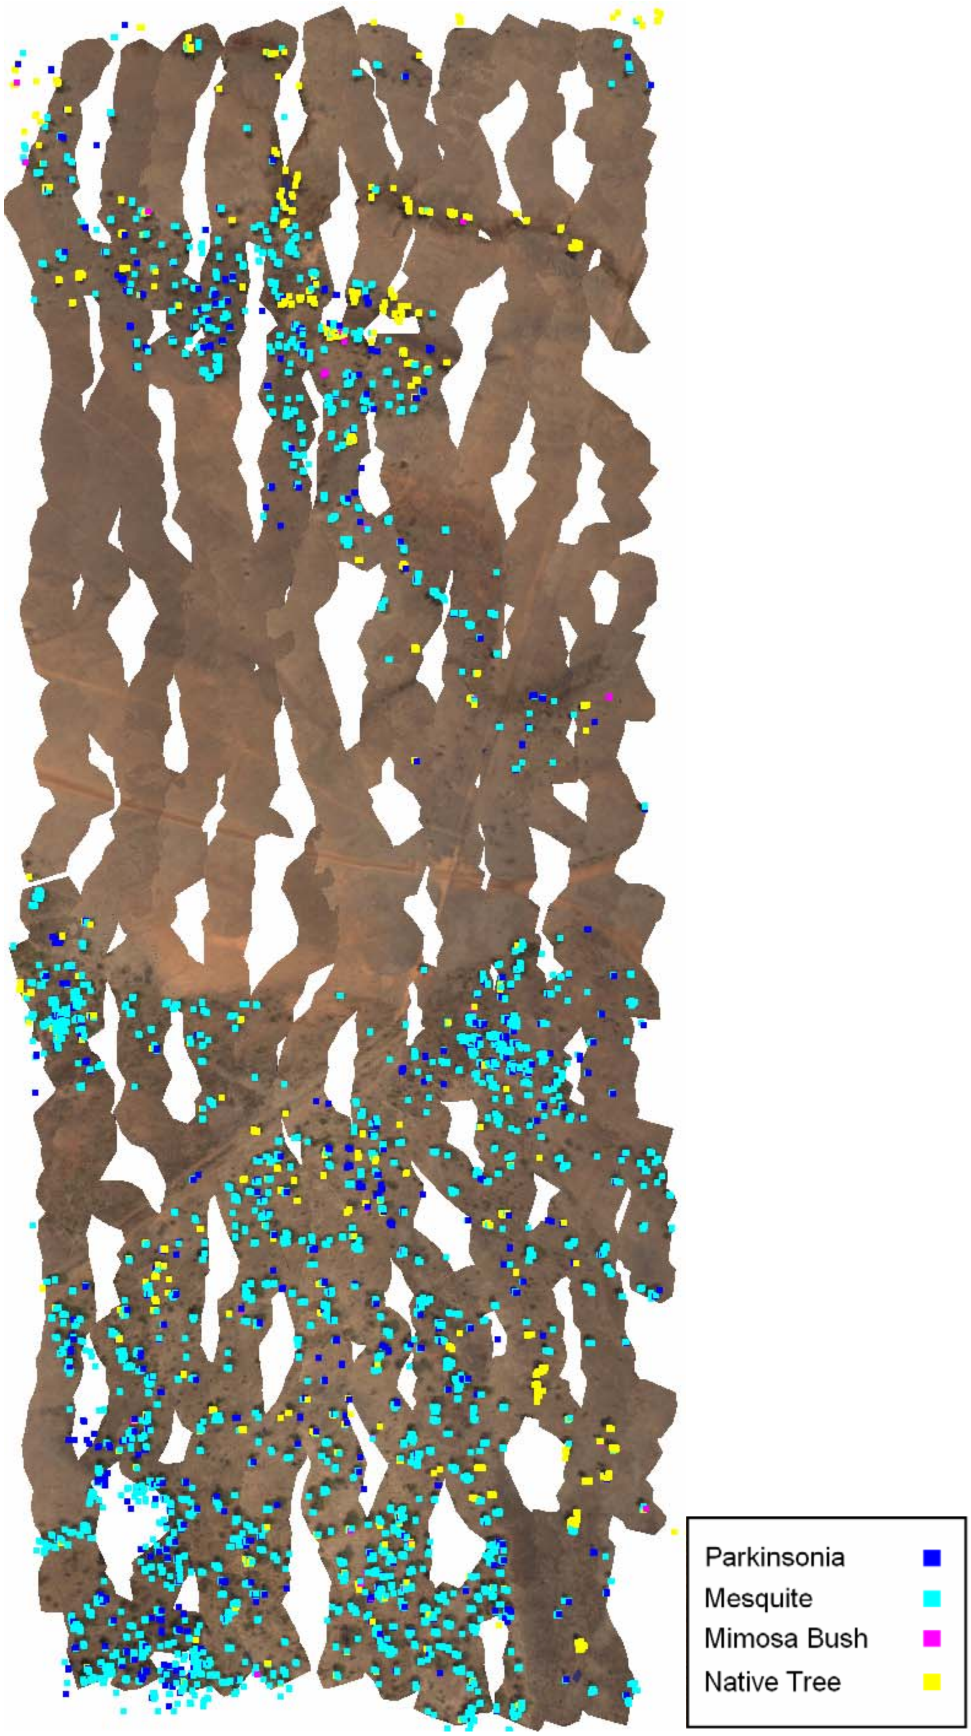

Figure A23 – Segment of constructed map of Williams Outstation from Flight 27 including classified tree crowns

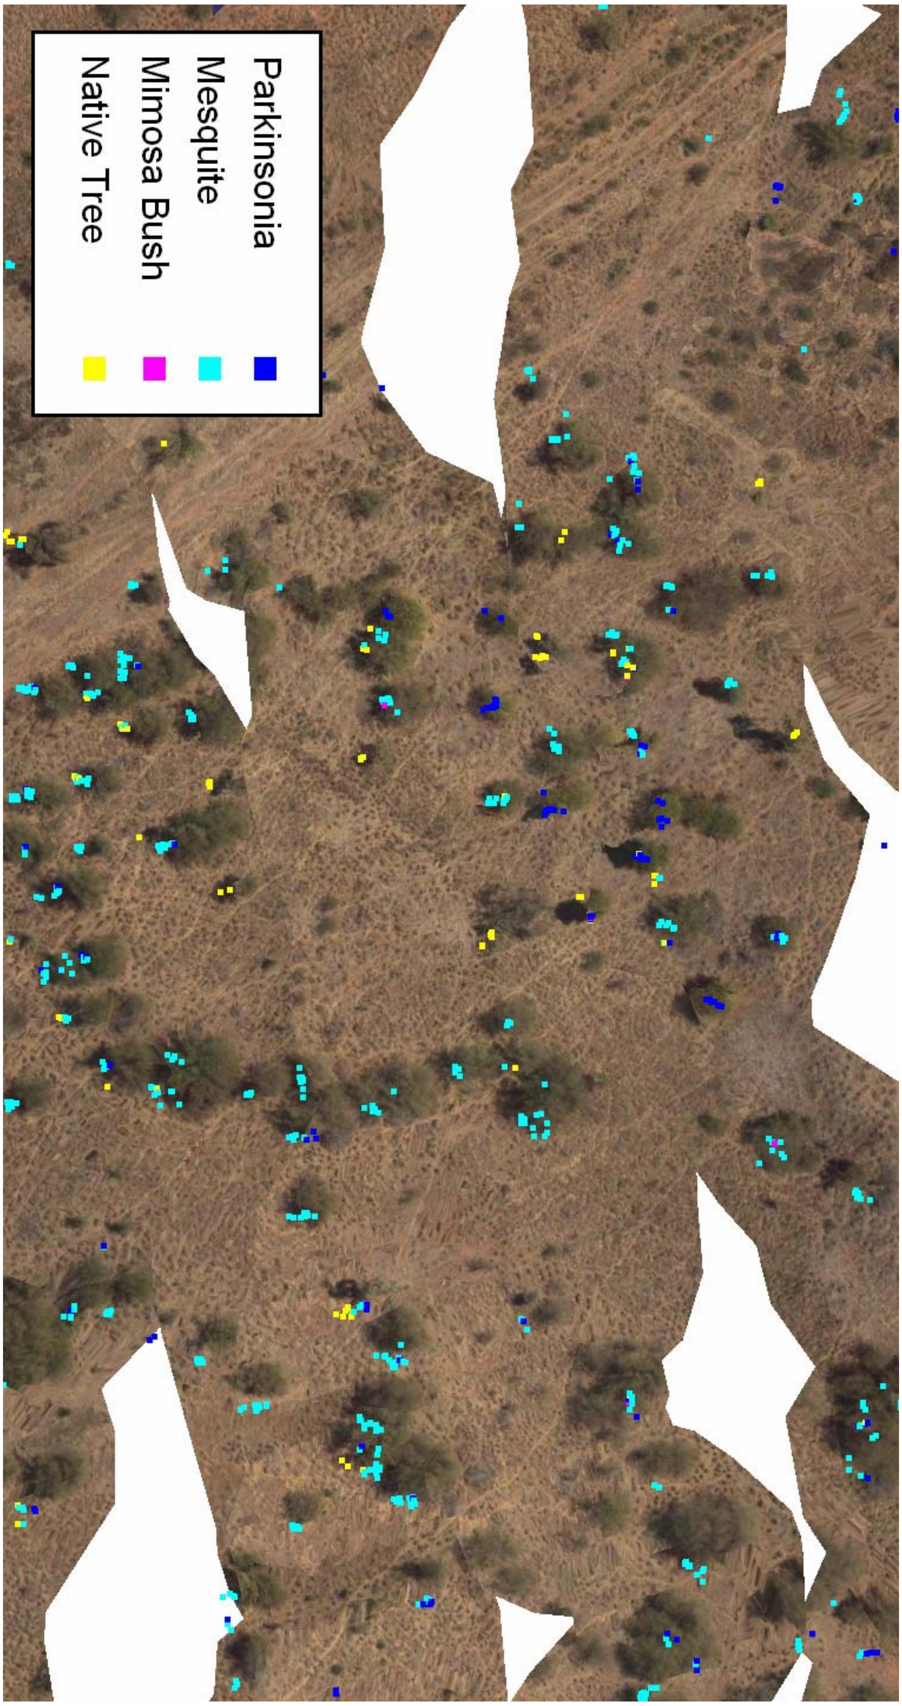

Figure A24 – Segment of constructed map of Williams Outstation from Flight 27 including classified tree crowns (zoomed-in)

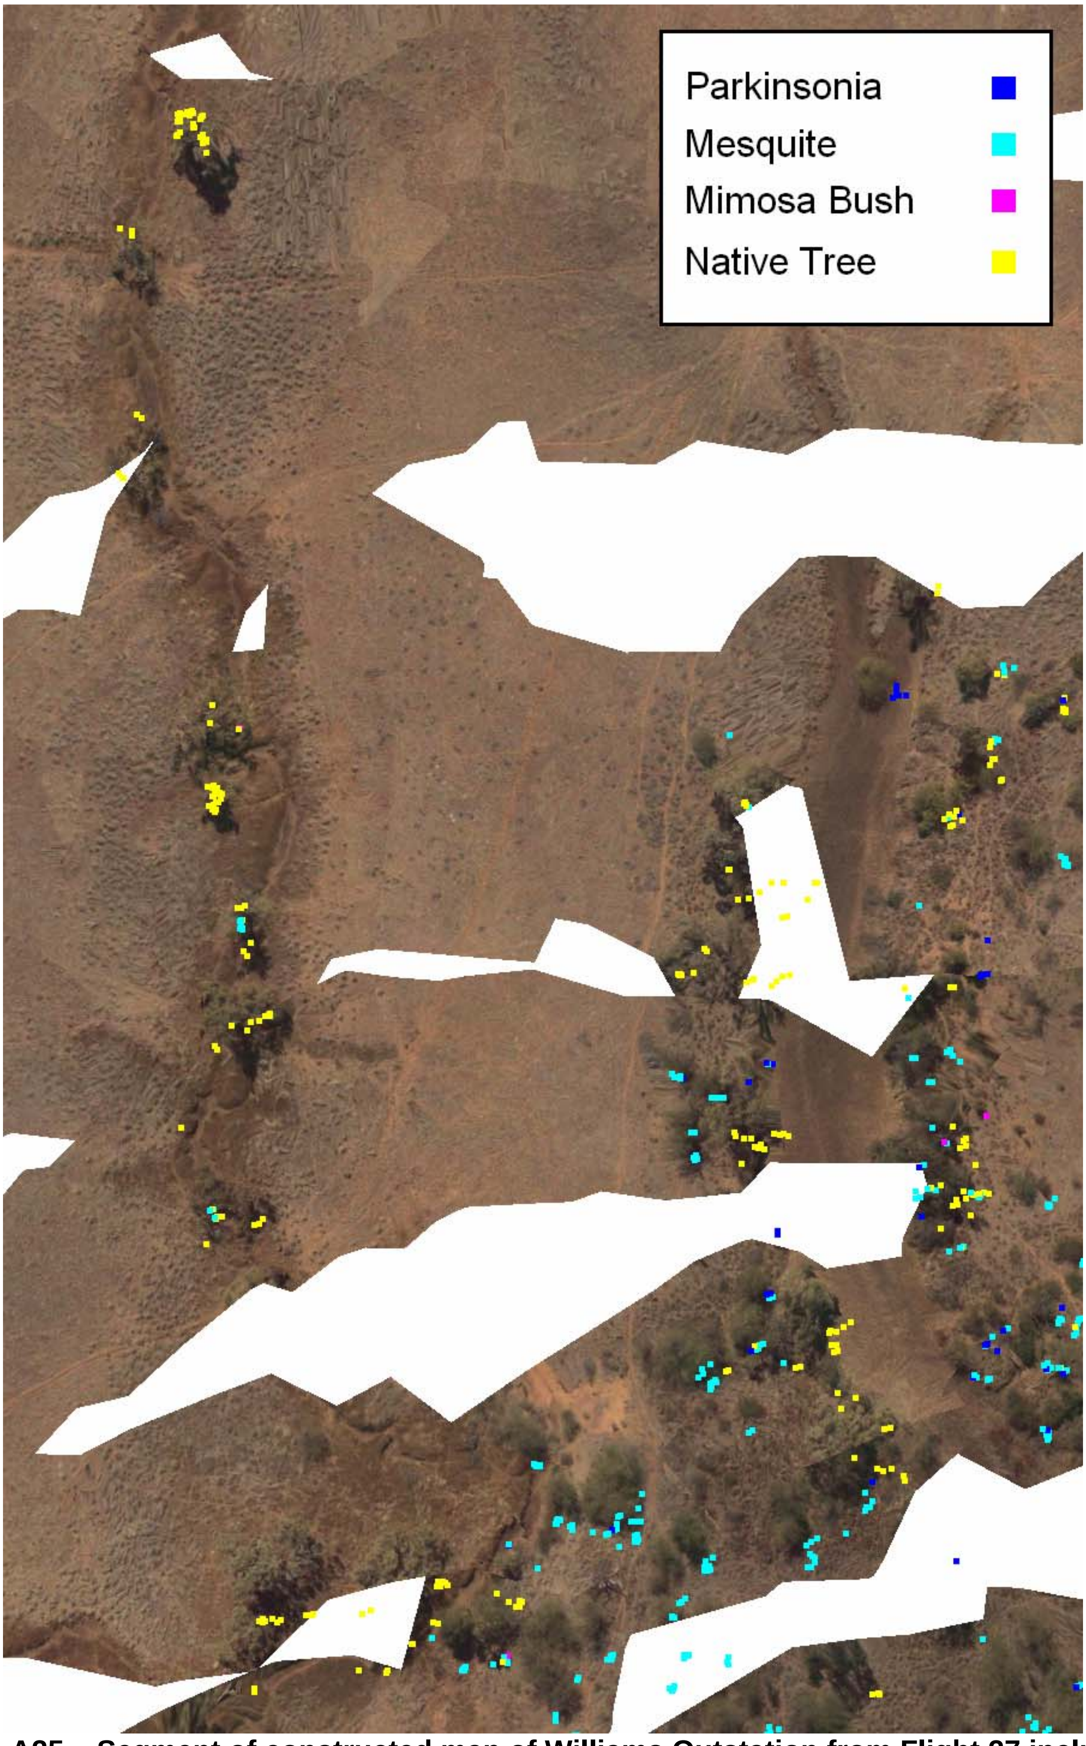

Figure A25 – Segment of constructed map of Williams Outstation from Flight 27 including classified tree crowns (zoomed-in)