Australian Government
Rural Industries Research and Development Corporation

LEADING THE SEARCH FOR WEED SOLUTIONS

Detecting Wheel Cacti

by using unmanned aerial vehicles and innovative classification algorithms

RIRDC www.rirdc.gov.au

Australian Government

Rural Industries Research and Development Corporation

LEADING THE SEARCH FOR WEED SOLUTIONS

Detecting Wheel Cacti –

By using unmanned aerial vehicles and innovative classification algorithms

by M Bryson and S Sukkarieh

May 2011

RIRDC Publication No 11/018
RIRDC Project No AWRC 08-04

ISBN 978-1-74254-205-8
ISSN 1440-6845

Using unmanned aerial vehicles and innovative classification algorithms in the detection of cacti
Publication No. 11/018
Project No. AWRC 08-04

Researcher contact details

Name: Dr Mitch Bryson
Address: Australian Centre for Field Robotics, University of Sydney

Phone: 02 9351 2186
Fax: 02 9351 7474
Email: m.Bryson@acfr.usyd.edu.au

Name: Assoc. Prof. Salah Sukkarieh
Address: Australian Centre for Field Robotics, University of Sydney

Phone: 02 9351 8154
Fax: 02 9351 7474
Email: salah@acfr.usyd.edu.au

In submitting this report, the researchers have agreed to RIRDC publishing this material in its edited form.

RIRDC contact details

Rural Industries Research and Development Corporation
Level 2, 15 National Circuit
BARTON ACT 2600

PO Box 4776
KINGSTON ACT 2604

Phone: 02 6271 4100
Fax: 02 6271 4199
Email: rirdc@rirdc.gov.au.
Web: http://www.rirdc.gov.au

Electronically published by RIRDC in May 2011
Print-on-demand by Union Offset Printing, Canberra at www.rirdc.gov.au
or phone 1300 634 313

Foreword

This project examined the role of unmanned aerial vehicles in detecting, classifying and mapping infestations of wheel cactus, *Opuntia robusta*, over large areas of rangelands in outback Australia.

Wheel cactus which is native and endemic to Mexico has now naturalised in South Australia, New South Wales and Victoria. It is often located in terrain which is difficult to access and monitoring and control by unmanned aerial vehicles and remote sensoring offers significant potential.

Trial flights by a fixed-wing unmanned aerial vehicle and an associated sensor payload took place near Oraparinna, in the Flinders Ranges of South Australia in September 2009 by staff and students from the Australian Centre for Field Robotics. Several flights were performed over an area of 1.2 by 2.0 kilometres, and data from the onboard sensor payload were logged and used to develop and test algorithms for the detection and mapping of cacti. In conjunction with the flight trails, a ground-based survey of cacti in the flight area was performed by ACFR staff and students in order to provide ground-truth and comparison data for validating the detection and mapping algorithms. This report presents details of the flight trials performed; details of the development of cactus detection, classification and mapping algorithms; and the results of these algorithms.

This project was funded in Phase 1 of the National Weeds and Productivity Research Program, which was managed by the Australian Government Department of Agriculture, Fisheries and Forestry (DAFF) from 2008 to 2010. The Rural Industries Research and Development Corporation (RIRDC) is now publishing the final reports of these projects.

Phase 2 of the Program, which is funded to 30 June 2012 by the Australian Government, is being managed by RIRDC with the goal of reducing the impact of invasive weeds on farm and forestry productivity as well as on biodiversity. RIRDC is commissioning some 50 projects that both extends on the research undertaken in Phase 1 and moves into new areas. These reports will be published in the second half of 2012.

This report is an addition to RIRDC's diverse range of over 2000 research publications which can be viewed and freely downloaded from our website www.rirdc.gov.au. Information on the Weeds Program is available online at www.rirdc.gov.au/weeds

Most of RIRDC's publications are available for viewing, free downloading or purchasing online at www.rirdc.gov.au. Purchases can also be made by phoning 1300 634 313.

Craig Burns
Managing Director
Rural Industries Research and Development Corporation

Contents

Tables

Figures

Executive Summary

What the report is about

This project examined the role of unmanned aerial vehicles in detecting, classifying and mapping infestations of wheel cactus, *Opuntia robusta*, over large areas of rangelands in outback Australia.

This report presents details of the flight trials performed; details of the development of cactus detection, classification and mapping algorithms; and the results of these algorithms.

Methods used

In the months leading up to September 2009 the focus was on preparing a fixed-wing unmanned aerial vehicle and an associated sensor payload—consisting of a terrain-facing monocular vision camera and onboard navigation sensors—for data collection over rangelands populated by cacti.

Trial flights near Oraparinna, in the Flinders Ranges of South Australia, were performed in September 2009 by staff and students from the Australian Centre for Field Robotics. Several flights were performed over an area of 1.2 by 2.0 kilometres, and data from the onboard sensor payload were logged and used to develop and test algorithms for the detection and mapping of cacti.

In conjunction with the flight trails, a ground-based survey of cacti in the flight area was performed by ACFR staff and students in order to provide ground-truth and comparison data for validating the detection and mapping algorithms.

Results/key findings

Results from the project demonstrate the applicability and capability of UAV systems using cheap sensors (such as a colour vision camera) for identifying and mapping cacti in remote locations. The success of the performance of the classification scheme was, however, only limited: the scheme demonstrated a good ability to detect cacti but only while providing a fairly large number of false detections.

Nevertheless, such a system in its current form could, for example, provide valuable information for a weed expert, acting as a first stage data-processing tool for identifying possible cactus locations but requiring the user to analyse and remove false detections.

Recommendations

Several avenues of future work for improving the system present themselves.

The first would be to examine vision-based feature descriptors that can better capture the visually distinguishing characteristics of the cacti. Apart from the colour and texture properties already exploited in this project, other possible features might be a type of shape detection—that is, a circular or elliptical Hough transform—in order to account for the 'wheel-like' leaves of the cactus.

A second avenue might be to vary imaging characteristics such as the capture resolution of the imagery data or to make observations of cacti from a more horizontal perspective by, for example, using a hovering vehicle to hover lower and closer to the ground. Wheel cacti are more naturally distinguishable from the horizontal perspective and so might be more easily classified if a second, lower flying vehicle were to be used to provide extra imagery.

Project activities

Preparation of the unmanned aerial vehicle platform and sensor payload

Efforts in the months leading up to the flight trials in September 2009 focused on ensuring that the J3 Cub fixed-wing UAV system (see Figures A.1 and A.2) and sensor payload were ready for collecting data in the target area of Oraparinna, in the Flinders Ranges of South Australia. A fixed-wing platform was chosen over a hovering platform because of the former's ability to cover most of the selected 1.2 by 2.0–kilometre area in a single flight. Efforts then focused on the design of the imaging sensor and the UAV's planned flight path given the following requirements for data collection:

- to provide imagery at a resolution of at least 3 centimetres/pixel on the ground or otherwise to maximise the resolution while operating at a safe altitude for the UAV

- to build three-dimensional spatial maps of detected objects in the imagery with a spatial accuracy of 5 metres (2σ) and to maximise coverage of the imagery while meeting the requirements for image resolution.

The image resolution requirements were determined primarily by the need to detect, identify and classify wheel cacti: existing aerial imagery of wheel cacti was unavailable, so the resolution requirements were estimated on the basis of required resolutions for identifying other vegetation of similar size, assuming an average cactus size of 1 square metre. Existing cactus ground survey data provided by Dylan Koener, the Oraparinna National Park ranger, were used for planning a flight path that would maximise the number of cacti in the expected coverage of the imagery (in order to maximise the number of aerial image training examples of cacti, as required for classification). The onboard inertial measuring unit GPS receiver provided information about the pose of the UAV, and this was used along with the camera information to build a spatial map of objects identified in the image data; these sensors allowed for an expected ground positioning accuracy of at least 5 metres.

Higher image resolution called for a lower operating altitude for the UAV and a narrower field-of-view camera lens, whereas maximising area coverage—and thus the number of cacti spotted—called for a higher operating altitude and a wider field-of-view lens. The trade-off was set by first minimising the operating altitude of the UAV to a safe height in view of the terrain of the flight area (200 metres above the lowest terrain points and 100 metres above the highest). Second, the lens field of view and flight pattern were set such that an area of 1.2 by 2 kilometres would be covered in a single one-hour flight, with overlap in the imagery so that a point on the ground had about a 90 per cent chance of being observed in a single flight (under errors in the coverage resulting from wind disturbance). It was planned there would be three flights over the same area in a single day in order to collect images at different sun illumination angles (to study the effect on the detectability of cacti) and to provide enough overlap of imagery to give about a 99 per cent chance that all ground points in the area would be covered by at least one image.

Table A.1 shows the final payload sensor specifications for the UAV system.

Flight trials

During September 2009 the Australian Centre for Field Robotics carried out field trials outside Oraparinna. Members of the ACFR team were involved in various activities—for example, managing and operating the platform, operating the ground station during flight, managing and operating the sensor payload box, processing data after flights, and performing ground-based manned surveys of cactus locations in the flight area to ground-truth and provide comparison data.

Four flights were made during the trials. The purpose of the first flight (referred to as Flight 0) was to test safe flying altitudes for the UAV. Because of the variations in the terrain of the flight area and the low elevation of the ground station site, the flight was used to test whether wireless radio communication between the UAV and the ground station would be lost as the UAV flew behind hills. The UAV flew at various altitudes, and it was found that radio communication was maintained for any altitude greater than 100 metres above the highest points in the terrain and 200 metres above the lowest. Three flights (referred to as Flights 1, 2, and 3) were performed on the following day, each lasting an hour and imaging the same 1.2 by 2–kilometre area. Table 1 provides details of the four flights.

Table 1 Operations during the September 2009 flight trials: a summary

Flight number	**Date**	**Area**	**Altitude (m)**
0	27 September	Oraparinna	Various
1	28 September	Oraparinna	200
2	28 September	Oraparinna	200
3	28 September	Oraparinna	200

Note: Flight 0 was a sensor payload test flight.

Figure A.3 shows the flight patterns for Flights 1, 2 and 3. A take-off and landing area at which the UAV ground station was located was designated just outside the flight area. In each flight the UAV flew along a fixed path that consisted of overlapping transects in order to provide a boxed coverage pattern. Figure A.4 shows the ground coverage of the camera footprint for each of the flights overlaid on the same figure. The footprint coverage is calculated on the basis of the navigation data (which provided information about the location of the UAV above the terrain), map information on terrain height (gathered from NASA SRTM data), and the knowledge of the camera field of view based on the camera lens used.

The ground survey in September 2009

In order to verify the ground survey data gathered before September 2009, staff and students of the Australian Centre for Field Robotics performed their own ground survey in the flight area. They found that several of the cacti in the data provided before September 2009 had been treated with herbicide and were no longer visible or had been dead so long that their remains were not visible from the air. The ACFR-run survey thus added several new cactus observations to the map. The survey was completed on foot in the area, a hand-held GPS receiver being used to record the locations of spotted cacti along with ground-based photographs and a description of the size and appearance of each cactus.

Figure A.5 shows the ground-surveyed cacti in the flight area. Thirty-eight cacti were located, along with other reference points and landmarks surrounding those specimens. The GPS coordinates from the ground survey allowed for isolation of camera frames of the aerial images that were expected to contain an image of the cacti. This was performed in conjunction with the self-navigation data collected by the UAV: combining the data allowed us to determine in which images a given cactus was visible and where in the image (with an accuracy of about +/– 100 pixels). In this way we could gather ground-truth aerial image patches of cacti to be used in training image detection algorithms to find cacti.

Figures A.6 to A.10 show examples of ground-based photographs and the corresponding aerial images of different cacti in the survey. Figure A.6 shows a cactus that the park ranger had treated with herbicide. In contrast, Figure A.7 shows two untreated cacti; these plants were on a cliff edge in reasonably inaccessible terrain and are fairly easy to discern from the ground-based, sideways-looking photograph but, because of their profile, much more difficult to discern in the bird’s-eye-view aerial image. Figures A.8 and A.9 show medium-sized cacti (about a cubic metre) in an area of sparse vegetation; the clear area around the cacti allows their shadows to be visible in the aerial imagery, a possible cue that could be used to help in their identification. Figure A.10 shows a very small cactus (about 30 cubic centimetres) that is extremely difficult to distinguish in the aerial imagery.

Development of cactus detection and mapping algorithms

Aerial image patches corresponding to the location of cacti identified in the ground survey were extracted from the camera data and used to design a cactus detection and classification scheme using aerial imagery. Figure A.11 shows image patches for different cacti identified in imagery data using ground survey and aircraft navigation information. A binary classifier was developed that would take as an input a small 64 x 48 patch of image data and return a classification of the image patch as corresponding to either a ‘cactus’ class or an ‘other’ class. Once the classifier was developed, each image from the aerial footage could be examined by breaking the image down into a grid of image patches and applying the classification to the entire set in the image.

Image patch vision features

Instead of classifying the raw pixel information in each image patch, image processing functions were used to build a vision feature vector—a list of parameters that described the colour, texture and shape of an image patch—that was used as the input to the binary classifier. The feature vector consisted of descriptors for the average colour information in the image patch (three values—one for each colour channel), as well as the responses from a Laplacian pyramid decomposition (Simoncelli & Freeman 1995) of the image patch in each of the three colour channels (five pyramid levels for each colour channel), which captures information about the roughness or smoothness of the texture information in the image patch at different scales (see Figure A.13). This resulted in a total of 18 parameters in the feature vector associated with an image patch.

Binary classifier training

The binary classifier was constructed by means of a learning procedure using manually extracted image patch training examples of both cactus and non-cactus image patches from the aerial imagery. Figure A.12 shows training image patches of non-cactus objects—other trees, bare ground, rocks, and so on—that were used in conjunction with the image patches of identified cacti in Figure A.11. The training procedure used a Logitboost algorithm (Friedman et al. 2000) that automatically assigns a hierarchy of simple decision rules on the basis of features identified in an image patch (known as ‘weak learners’) and so minimises classification errors given the training data set. We used the algorithm because of its ease of ‘tuning’ (with only one design parameter—the number of weak learners used in the hierarchy) and its low computational complexity for learning and inference. Logitboost is a ‘boosting’ algorithm that uses the outputs of multiple weak, simple learning algorithms to make stronger inferences than the weak algorithms could achieve individually.

Cactus mapping

Once the cactus detection and classification algorithms have been used to identify cacti in each image of the data collected by the UAV, the corresponding location of each cactus in the image needs to be mapped into the environment by using information about the UAV’s position and orientation, as provided by the onboard navigation system. This was done in a two-step process.

The first step involved using collected sensor information from the vision camera and onboard navigation sensors (inertial measuring unit and GPS) to estimate the position and orientation of the UAV at each camera frame and to produce a geo-referenced point feature map of the mapped terrain. The estimation was done using a batch simultaneous localisation and mapping, or SLAM, algorithm

(Bryson et al. forthcoming) that incorporated all the available sensor information to produce a statistically least-square error optimal map. Figure A.14 shows the relationship between the position of the UAV, the position of a cactus in the environment, and the relative camera observation made of the cactus from the UAV.

The second step involved using the UAV information on position and orientation, the terrain map and pixel location of identified cacti in the vision data to triangulate the three-dimensional position of each cactus in the environment. This was done by computing the line-of-sight intersection of the cactus observation with the estimated terrain surface.

Cactus detection and mapping results

In order to evaluate the classification algorithm, a 10 per cent out cross-validation was performed. For this we used a collection of 50 image patches of identified cacti and 500 image patches of other parts of the terrain identified as not corresponding to cacti. The larger number of ‘other’ image patches compared with cactus image patches was chosen to reflect the fact that cacti were observed far less frequently in the environment than in other terrain; this would provide a more realistic performance evaluation of the algorithms. The cross-validation was done by choosing one of the image patches, removing this image patch along with another 10 per cent of the image patches (randomly selected) and training the Logitboost binary classifier using the remaining image patches. The trained classifier was then used to classify the originally removed image patch, and the result was recorded as either a correct or an incorrect classification. This process was repeated by cycling through each of the image patches (both cactus and other patches).

The classification performance was measured using three main metrics. The first, accuracy, is defined as the percentage of correct classifications of both cacti and other objects. The second, precision, has a separate value for each of the two classes and is defined as the percentage of correct classifications of that class compared with the total number of objects examined in that class. The third metric, recall, which also has a separate value for each of the two classes, is defined as the ratio of correct classifications to the sum of the correct classifications and incorrect classifications of the other class (that is, missed classifications of the original class).

In order to tune the parameter in the classifier (that is, the number of weak learners used) the cross-validation procedure was repeated for several different values of weak learners used during training. Figure A.15 shows the classification performance values as a function of the number of weak learners used for classification. It is expected that as the number of weak learners is increased the possible complexity of the classifier, and thus the flexibility in relation to different types of input vectors, should increase. It is also expected that as the number of weak learners is increased the classifier will become more and more specific to the actual input data used for training and thus might not be sufficiently general or robust to identify new types of image data. The trade-off is achieved by increasing the number of weak learners to the point where no large performance gains are evident. We found that a value of 50 weak learners was appropriate given the observed data and feature vector, so this value was used in the final classifier.

Tables 2 and 3 show the final performance values for the cactus classifier. The confusion matrix shown in Table 2 shows the number of cases for the four possible outcomes—correctly identified as cactus, incorrectly identified as cactus, correctly identified as other, and incorrectly identified as other. Table 3 shows the performance values of accuracy, precision and recall computed from the cross-validation results. The results suggest a high degree of precision in the identification of cacti; in other words, the classifier rarely misses a cactus by erroneously classifying it as another object in the image. The results also suggest a poor recall in the identification of cacti; in other words, even though most cacti are identified, a fairly high number of other image segments are also identified as being cacti when they are not.

Table 2 Confusion matrix of the classification results using a 10 per cent out cross-validation

		Actual	
Classified		**Cacti**	**Others**
	Cactus	46	82
	Other	4	418

Note: Table shows the number of image patches correctly identified by the classifier (for cacti and others, diagonal terms) and the number of image patches incorrectly identified (for cacti and others, off-diagonal terms).

Table 3 Precision, recall and accuracy values for the classifier

Value	**Cactus (%)**	**Others (%)**
Precision	92	84
Recall	36	99
Overall accuracy (%)	84	

Note: Accuracy is the total percentage of correct classifications. For a given class (that is, cacti or others) precision is the number of correct classifications of this class divided by the total number in the class, and recall is the number of correct classifications divided by the sum of the correct classifications and the incorrect classifications of other classes (that is, erroneous classifications).

Figure A.16 shows examples of cacti and the ‘other’ classification in selected vision frames known to contain cacti. Overlaid in the images are the locations of known cacti. It is evident that all the cacti in the frames have been identified correctly but that a large number of other non-cactus objects in the frame have also been labelled as cacti, supporting the precision and recall results seen in Table 3.

Project outcomes and conclusions

The project studied the use of low-flying UAVs for detection and mapping of wheel cacti over large areas of rangelands. The UAV deployed carried a sensor payload consisting of a monocular colour vision camera and navigation sensors, which allowed the UAV to detect cacti in the terrain over which it flew while mapping the three-dimensional location of detected cacti on the basis of the identified pixel position in the camera and the onboard navigation information. Flight trials were performed over rangelands near Oraparinna in the Flinders Ranges of South Australia, and logged sensor data were used to aid in the development of and to test classification and mapping algorithms.

Results from the project demonstrate the applicability and capability of UAV systems using cheap sensors (such as a colour vision camera) for identifying and mapping cacti in remote locations. The success of the performance of the classification scheme was, however, only limited: the scheme demonstrated a good ability to detect cacti but only while providing a fairly large number of false detections. Nevertheless, such a system in its current form could, for example, provide valuable information for a weed expert, acting as a first stage data-processing tool for identifying possible cactus locations but requiring the user to analyse and remove false detections. From the aerial photographs provided in Figures A.6 to A.11 it can be seen that aerial identification of cacti in this context is inherently difficult, most cacti being very difficult to spot (even by eye) and there being very few distinguishing characteristics.

Several avenues of future work for improving the system present themselves. The first would be to examine vision-based feature descriptors that can better capture the visually distinguishing characteristics of the cacti. Apart from the colour and texture properties already exploited in this project, other possible features might be a type of shape detection—that is, a circular or elliptical Hough transform (Rizon et al. 2005)—in order to account for the ‘wheel-like’ leaves of the cactus. A second avenue might be to vary imaging characteristics such as the capture resolution of the imagery data or to make observations of cacti from a more horizontal perspective by, for example, using a hovering vehicle to hover lower and closer to the ground. Wheel cacti are more naturally distinguishable from the horizontal perspective (for example, see Figure A.7) and so might be more easily classified if a second, lower flying vehicle were to be used to provide extra imagery.

References

Bryson, M, Reid, A, Ramos, F & Sukkarieh, S (forthcoming), ‘Airborne vision-based mapping and classification of large farmland environments’, *Journal of Field Robotics* (accepted January 2010).

Friedman, J, Hastie, T & Tibshirani, R 2000, ‘Additive logistic regression: a statistical view of boosting’, *Annals of Statistics*, vol. 28, no. 2, pp. 337–407.

Rizon, M, Yazid, H, Saad, P, Shakaff, A, Saad, R, Sugisaka, M et al. 2005, ‘Object detection using circular Hough transform’, *American Journal of Applied Sciences*, vol. 2, no. 12.

Simoncelli, E & Freeman, W 1995, ‘The steerable pyramid: a flexible architecture for multi-scale derivative computation’, *2nd IEEE International Conference on Image Processing*, vol. 3, pp. 444–7.

Appendix A Supporting information

Flight platform and sensor payload system

Figure A.1 J3 Cub UAV system in flight over Oraparinna, Flinders Ranges, South Australia, September 2009

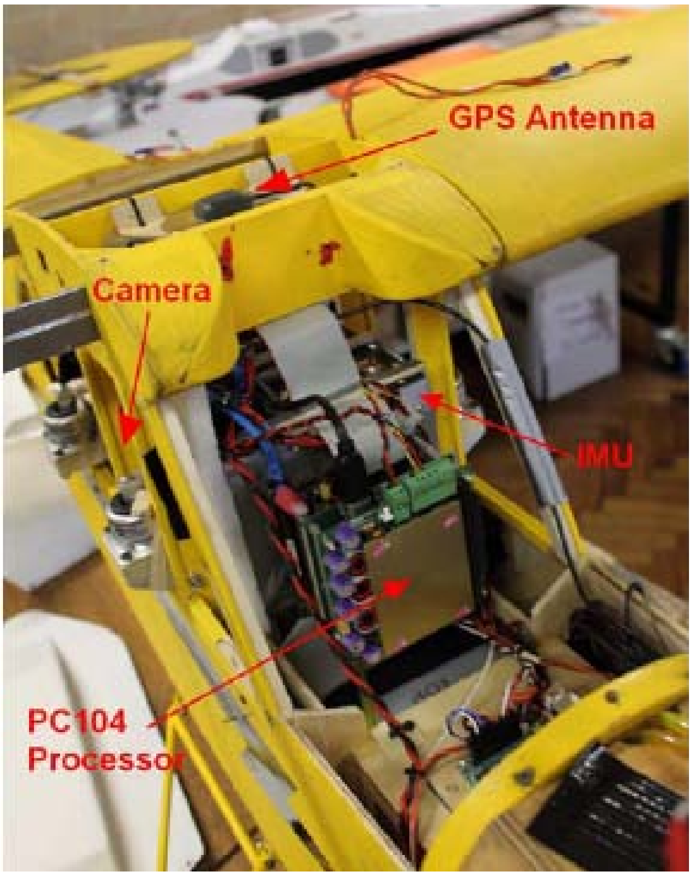

Figure A.2 Internal view of the J3 Cub UAV system, showing payload box incorporating the inertial measuring unit, the GPS receiver and antenna, the vision-stack processor and the camera system

Table A.1 Sensor payload specifications

Vision camera	**Hitachi HV-F31**	**Inertial measuring unit**	**Honeywell HG1900**
Sampling rate	3.75Hz	Sample rate	600 Hz, pre-processed to 100 Hz
Field of view	10.9° x 8.2°	Accel. noise (1σ)	0.05 m/s^2
Resolution	1024 x 768 pix	Gyro noise (1σ)	0.05°/s
Angular resolution	0.0106°	Accel. bias (1σ)	0.05 m/s^2
Ground resolution	2.8 cm/pix @ 200 m	Gyro bias (1σ)	0.05°/s
Ground footprint	38 x 30 m @ 200 m		

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

Note: The sensor payload consists of an inertial measuring unit, a GPS receiver and a downwards-mounted colour monocular camera.

Oraparinna flight trials

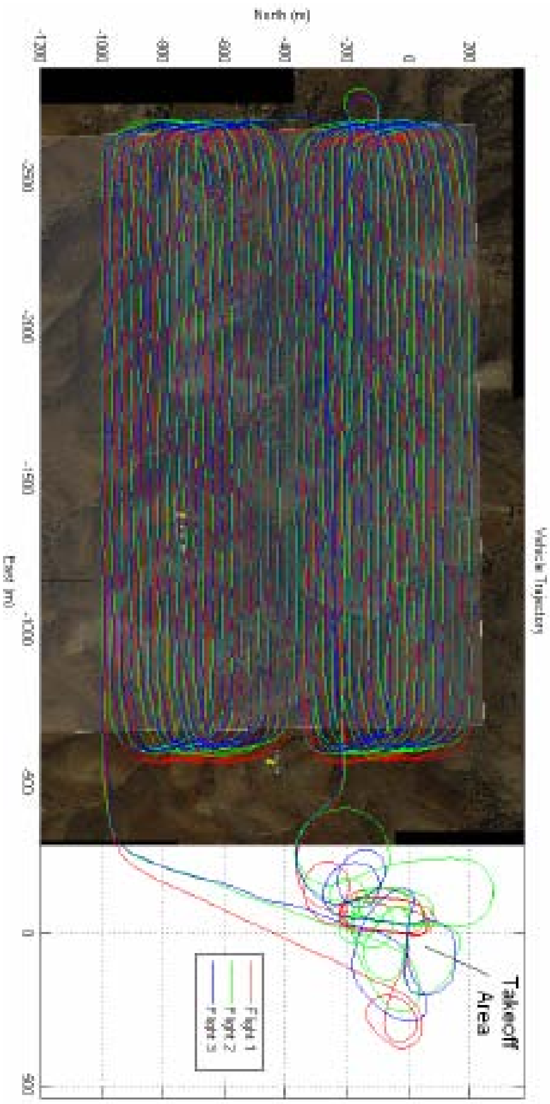

Note: Shown underneath the flight paths is low-resolution imagery of the area available from Google Earth.

Figure A.3 Flight paths for Flights 1, 2 and 3 at Oraparinna

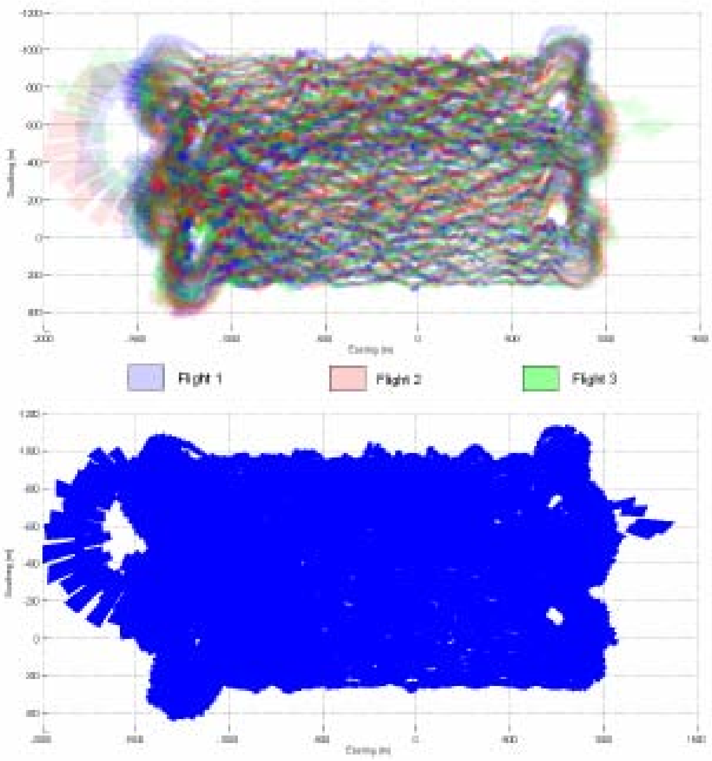

Note: The top figure shows coverage footprints for the three different flights where darker areas indicate sections of the ground with coverage in multiple images. The bottom figure shows areas of coverage where at least one image is available between the three flights.

Figure A.4 Imagery coverage patterns for the three flights performed at Oraparinna

Oraparinna ground-truth cactus survey and aerial photos

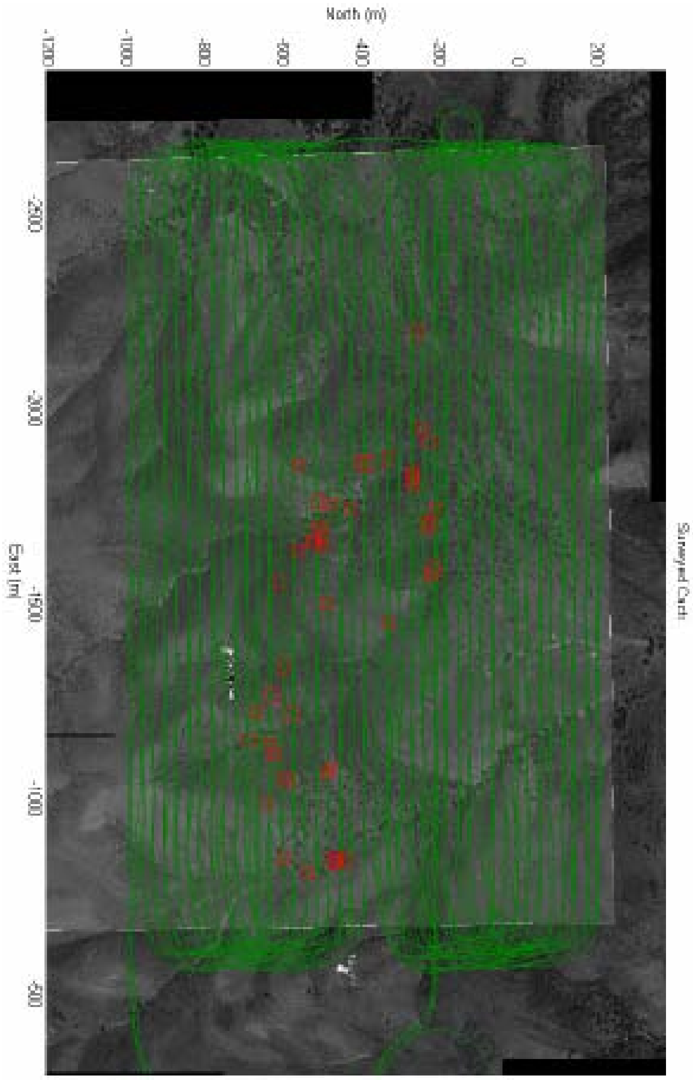

Note: Cactus locations in red; UAV flight path in green.

Figure A.5 Ground survey cactus locations with UAV flight path information overlaid

Note: Aerial photo location provided by UAV navigation system and ground-surveyed position data.

Figure A.6 Ground-based and aerial photography of wheel cactus #1, a herbicide-treated cactus

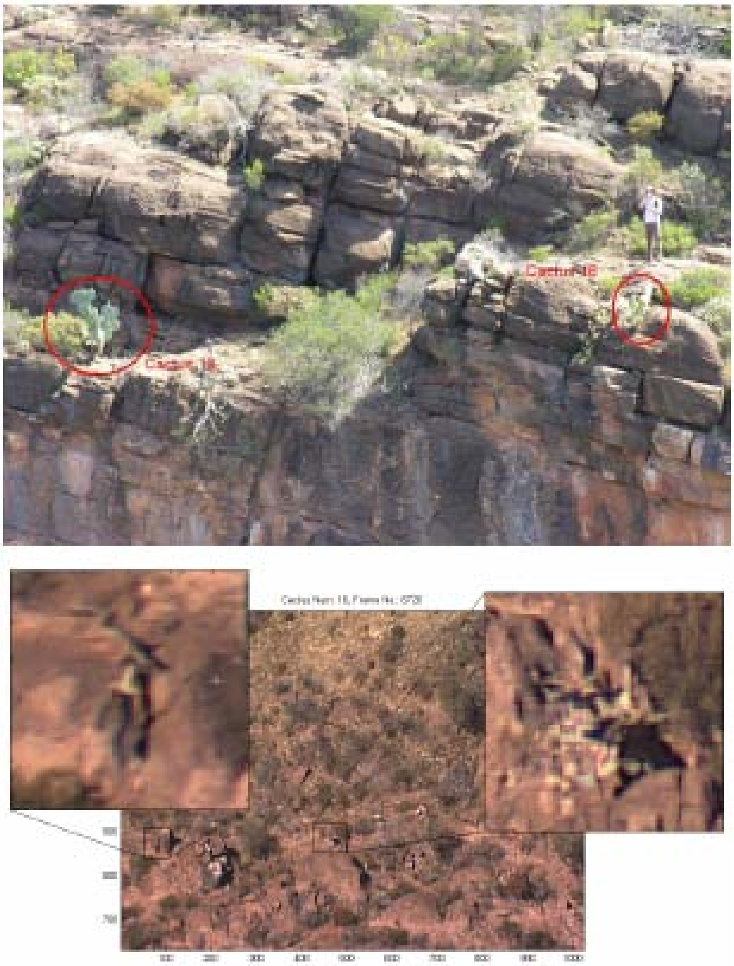

Note: Aerial photo location provided by UAV navigation system and ground-surveyed position data.

Figure A.7 Ground-based and aerial photography of wheel cacti #16 and #18, two untreated cacti located on a cliff ledge

Note: Aerial photo location provided by UAV navigation system and ground-surveyed position data.

Figure A.8 Ground-based and aerial photography of wheel cactus #22

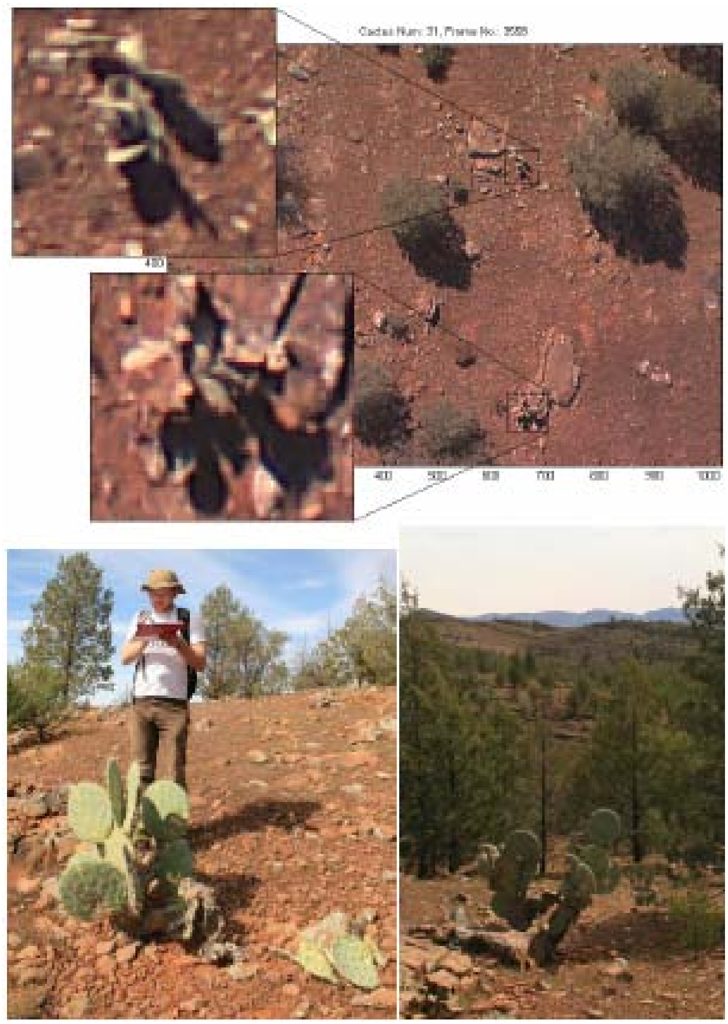

Note: Aerial photo location provided by UAV navigation system and ground-surveyed position data.

Figure A.9 Ground-based and aerial photography of wheel cacti #31 and #32, two medium- to large-sized cacti

Note: Aerial photo location provided by UAV navigation system and ground-surveyed position data.

Figure A.10 Ground-based and aerial photography of a very small (about 20 cubic centimetres) wheel cactus

Note: The image patch is extracted from the aerial footage using information about the position of the cactus from the survey and information about the position and orientation of the UAV from navigation information.

Figure A.11 Collected image patches (64 x 48 pixels) of different cacti located during the ground survey

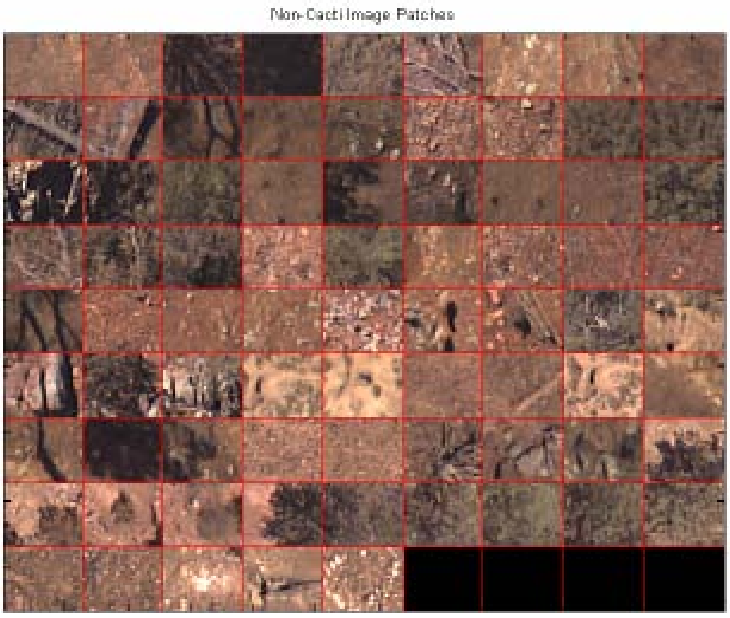

Note: These image patches were used as training counter-examples for cactus image detection algorithms.

Figure A.12 Collected image patches (64 x 48 pixels) of different features in the aerial imagery (non-cactus)

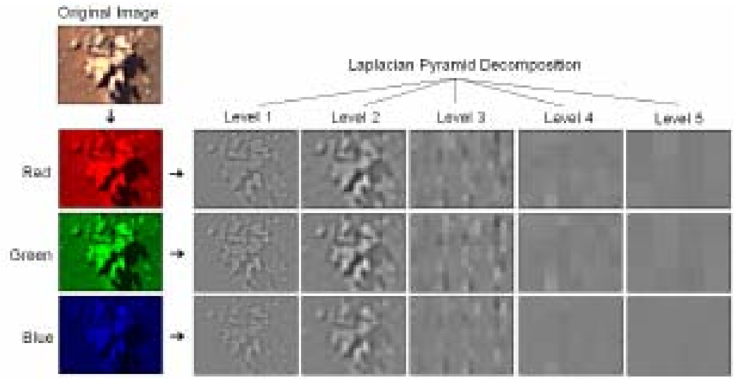

Note: Eighteen parameters make up the feature descriptor vector for an image patch consisting of the average colour response in each of the red, green and blue channels of the image (three parameters) and the average Laplacian pyramid-level response in each colour channel, across five pyramid levels.

Figure A.13 Image patch feature descriptors

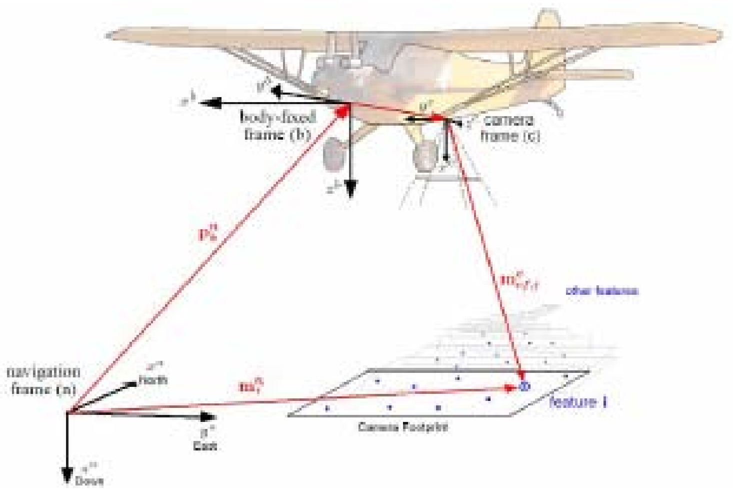

Figure A.14 Relationships between UAV position and orientation, feature (cactus) position and observation by the on-board camera, all used for mapping the location of observed cacti

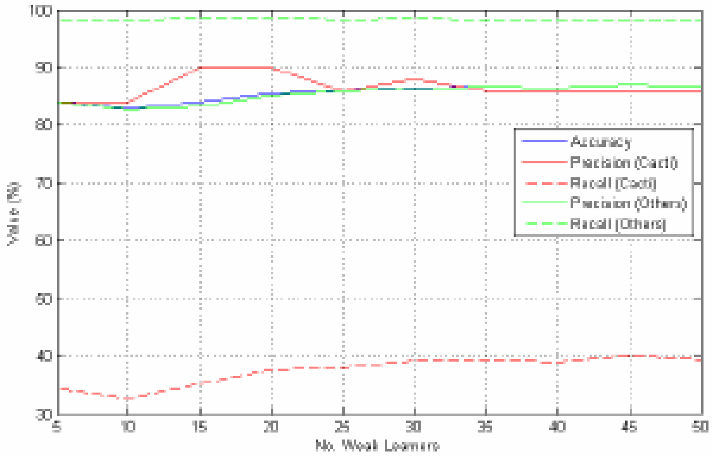

Note: The results show a limited performance gain from using up to 50 weak learners (over lower numbers) but almost no gain for a more complex classifier with high numbers.

Figure A.15 Classification results—accuracy, precision and recall—using a 10 per cent out cross-validation versus the number of weak learners used during classification

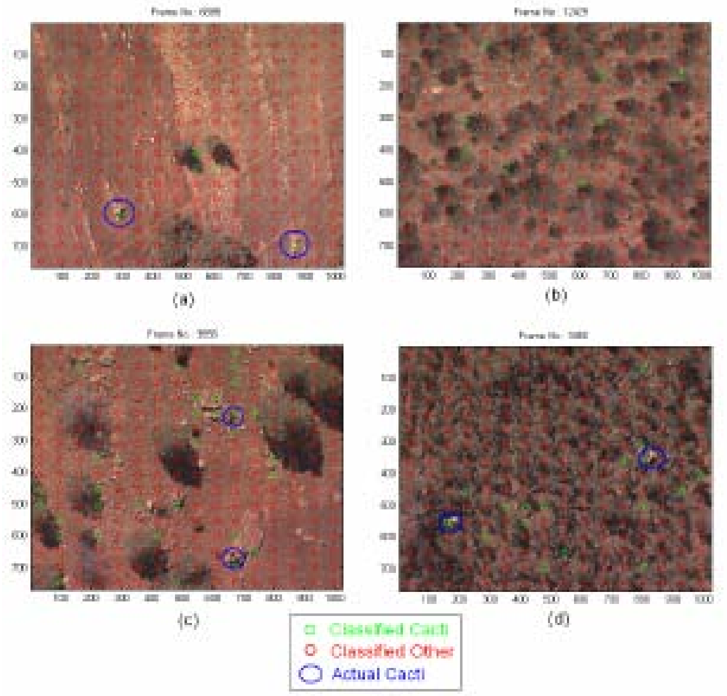

Note: The classifier detects all cacti within the image frame (high precision) but also erroneously classifies a large number of non-cactus objects as cacti (poor recall).

Figure A.16 Selection of classified image frames with overlay of actual cactus locations

Detecting Wheel Cacti – *by using unmanned aerial vehicles and innovative classification algorithms*

by m Bryson and S Sukkarieh

This project examined the role of unmanned aerial vehicles in detecting, classifying and mapping infestations of wheel cactus, *Opuntia robusta*, over large areas of rangelands in outback Australia. In the months leading up to September 2009 the focus was on preparing a fixed-wing unmanned aerial vehicle and an associated sensor payload—consisting of a terrain-facing monocular vision camera and onboard navigation sensors—for data collection over rangelands populated by cacti. Trial flights near Oraparinna, in the Flinders Ranges of South Australia, were performed in September 2009 by staff and students from the Australian Centre for Field Robotics. Several flights were performed over an area of 1.2 by 2.0 kilometres, and data from the onboard sensor payload were logged and used to develop and test algorithms for the detection and mapping of cacti. In conjunction with the flight trails, a ground-based survey of cacti in the flight area was performed by ACFR staff and students in order to provide ground-truth and comparison data for validating the detection and mapping algorithms.

This report presents details of the flight trials performed; details of the development of cactus detection, classification and mapping algorithms; and the results of these algorithms.

This project was funded in Phase 1 of the National Weeds and Productivity Research Program, which was managed by the Australian Government Department of Agriculture, Fisheries and Forestry (DAFF) from 2008 to 2010. The Rural Industries Research and Development Corporation (RIRDC) is now publishing the final reports of these projects.

Solutions to weeds in Australia require a long-term, integrated, multi-stakeholder and multi-disciplinary approach. RIRDC is seeking project applications that involve collaboration between stakeholder groups, and where possible, including external contributions both monetary and in-kind.

This report is an addition to RIRDC's diverse range of over 2000 research publications which can be viewed and freely downloaded from our website www.rirdc.gov.au.. Information on the Weeds Program is available online at www.rirdc.gov.au/weeds. Most of RIRDC's publications are available for viewing, free downloading or purchasing online at www.rirdc.gov.au. Purchases can also be made by phoning 1300 634 313.

Cover photos: Front cover: Wheel cactus.
Back cover photos are sourced from this report.

This publication can be viewed at our website—www.rirdc.gov.au.
All RIRDC books can be purchased from:

www.rirdc.gov.au

Contact RIRDC:
Level 2, 15 National Circuit
Barton ACT 2600

PO Box 4776
Kingston ACT 2604

Ph: 02 6271 4100
Fax: 02 6271 4199
Email: rirdc@rirdc.gov.au
web: www.rirdc.gov.au

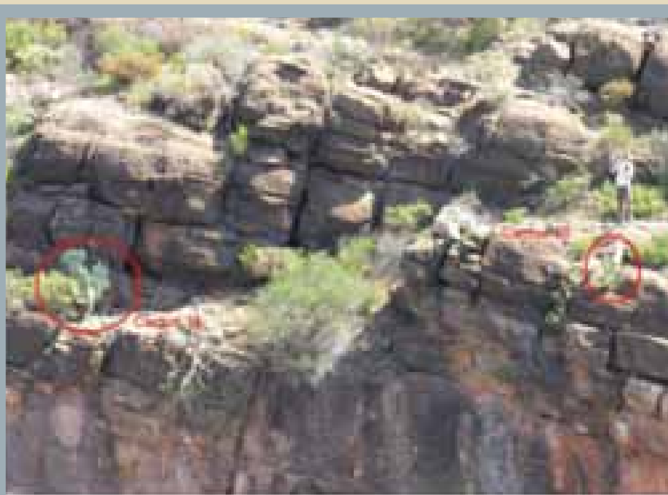

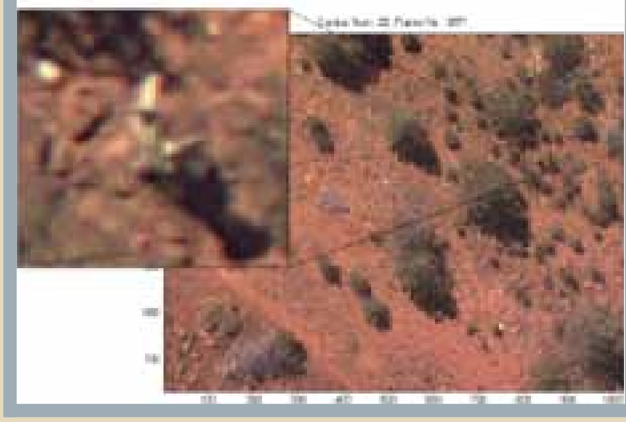

RIRDC www.rirdc.gov.au