Horticulture
Innovation
Australia

Final Report

Precision weed sensing for pyrethrum

Cheryl McCarthy
National Centre for Engineering in Agriculture

Project Number: PY12000

PY12000

This project has been funded by Horticulture Innovation Australia Limited with co-investment from Botanical Resource Australia and funds from the Australian Government.

Horticulture Innovation Australia Limited (Hort Innovation) makes no representations and expressly disclaims all warranties (to the extent permitted by law) about the accuracy, completeness, or currency of information in *Precision weed sensing for pyrethrum.*

Reliance on any information provided by Hort Innovation is entirely at your own risk. Hort Innovation is not responsible for, and will not be liable for, any loss, damage, claim, expense, cost (including legal costs) or other liability arising in any way (including from Hort Innovation or any other person's negligence or otherwise) from your use or non-use of *Precision weed sensing for pyrethrum*, or from reliance on information contained in the material or that Hort Innovation provides to you by any other means.

ISBN 0 7341 3899 7

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

Contents

Summary

USQ-NCEA has developed and evaluated a proof-of-concept precision weed sensing system for the pyrethrum industry to improve productivity and reduce the variable costs of production caused by weeds. The precision weed sensing system was based on machine vision, building on research by NCEA since 2007 in the sugar, grains and cotton industries. Machine vision combines the functionality of a camera, computer and software to programmatically identify information such as weeds or crop in images. The project was undertaken in collaboration with Botanical Resources Australia (BRA). The target audience for this report are pyrethrum industry stakeholders, following a confidentiality period of at least one year while the technologies that were developed in the project are commercialised.

A test-rig and data collection unit was developed to capture image data of pyrethrum under a range of weed conditions for development of a proof-of-concept image analysis algorithm. As approved in a project variation, a pre-production spot spray prototype was developed and evaluated that incorporated a precision spray delivery system for on-the-row and between-the-row spraying, a guidance system to maintain the implement over the pyrethrum rows, and a precision weed sensing system incorporating image analysis algorithms and machine vision equipment to recognise weeds and trigger spray solenoids. The technology recorded image and GPS data to build image maps of pyrethrum rows.

The proof of concept and pre-production prototype were evaluated on 60cm and 45cm row spacings to enable weed management between-the-row and on-the-row, which was being trialed by BRA. Different row spacings led to different requirements for image analysis development. The image analysis was found to be 98% successful in pyrethrum and weed discrimination with 1.2% overspray on pyrethrum for 60cm row spacings.

The project has delivered outputs as per the project plan, which were: single row test rig and proof-of-concept algorithm; field prototype, performance specifications and usage guidelines for the technology; final performance specification; and final report. Results were presented at industry forums and are planned to be published in scientific journals. Specific key outputs of the project have been:

- patentable algorithms and processing techniques that enables real-time operation in green from green image analysis in horticulture as well as wider agricultural systems;
- domain knowledge on commercial spray apparatus and its capability in the field; and
- real-time map generation for ‘data mining’ such as plant health, plant density and weed maps.

The outcomes of the project have been achieved which were to deliver performance specifications on wide row pyrethrum for the proof of concept and pre-production technology. Between row technology has been trialled and weed map data has been recorded which might be used to identify emerging weed problems and inform weed management strategies. Specific key outcomes from the project were:

- KPIs have been met, namely real-time operation of a pre-production prototype spot sprayer with 98% hit rate of weeds with 100mm diameter, traveling at speeds up to 8 km/h in 60cm pyrethrum row spacing.
- Ongoing commercialisation of the technology by USQ, SRA, HIAL and CRDC, which will produce a commercial ready sensor that can be bought from a commercial company and used in the pyrethrum industry, and identification of a potential commercial partner in John Deere.

Key recommendations for follow on research are: use on narrow rows, which is already underway in a separate project between USQ and BRA; application to the wider horticulture and agricultural industries; and data mining of recorded machine vision images for other precision agriculture uses.

The project has been pivotal in ongoing efforts to commercialise precision weed sensing technology and provide a commercial product that can purchased, used and maintained by farmers and contractors into the future. This will take the research to a reality and provide much needed new weed management tools to the pyrethrum and wider agricultural industries.

Keywords

Machine vision, image analysis, depth camera, real-time, green from green, row guidance, spot spray, spray nozzles, weed control, commercialisation

Introduction

Pyrethrum (*Chrysanthemum cinerariifolium* (Trev.) Vis.) is a perennial plant that produces daisies which can be harvested over successive years. Throughout the year pyrethrum experiences periods of slow to fast growth from one year's harvest to the next. The variation in the pyrethrum crop growth and canopy throughout this cycle creates a number of weed management challenges. When a pyrethrum crop is in its slow growing phase from harvest through to spring it is susceptible to increased competition and infestation from surrounding weeds. When weeds are not controlled effectively within a pyrethrum crop the weed seed bank can build up compounding weed control issues and competition in subsequent seasons. In cases where herbicide is applied and does not kill weeds effectively weed resistance to the herbicide applied can occur and alternative herbicide combinations are required to break the weeds resistance. When selective herbicides are applied certain weed groups that are not susceptible to the herbicide applied may persist and generate seed for the following season. As pyrethrum is harvested using a fleet of harvest equipment that is moved from field to field there is risk for weed seed caught in harvesting equipment to be distributed and identification and control of emerging weed varieties unpredictable.

To combat these challenges improved weed identification and control is required to inform and enhance weed management strategies. Machine vision is a viable methodology to identify weeds as it combines the functionality of camera, computer and software together in a conceptually similar manner to our eyes, brain and knowledge to artificially identify information such as weeds or crop in images. The artificial identification aspect of machine vision system would provide several benefits over the current hand-spot spraying and broadcast spraying methods adopted by the industry. A precision weed sensing system encompassing a machine vision system would have the capability to identify and spray individual weeds amongst the crop in a similar manner to hand-spot spraying while traveling at speeds comparable to broadcast spraying.

Prior to the commencement of this project the National Centre for Engineering in Agriculture has been involved in weed identification projects funded by Sugar Research and Development Corporation (SRDC) (now Sugar Research Australia) and Rural Industries Research and Development Corporation (RIRDC) for cotton and grains. The proposed Horticulture Australia Limited (HAL) project (now Horticulture Innovation Australia (HIA)) aimed to continue this research and enable demonstration, further refinement and full evaluation of the proof-of-concept precision weed sensing system in pyrethrum. Existing commercially available spot spray systems (i.e. Weedit, WeedSeeker) are not suitable for identifying green weeds amongst green crop and can only detect weeds that are present over stubble or soil backgrounds. Garfords Robocrop provides an on-the-row mechanical weed-control solution in situations where plant is the dominant feature in the image, the plant and row spacing's are consistent and consecutive plants on-the-row are not touching. This HIA project developed a precision weed sensing system to expand the pyrethrum industries weed control methods and overcome limitations with existing spot spray technologies that are unable to differentiate between plant species and detect green weeds from green pyrethrum plants that are not evenly spaced and are touching one another.

A proof-of-concept spot-spray prototype-implement was built to spray weeds from within the pyrethrum crop that consisted of a precision spray delivery system for on-the-row spraying, spray shields for between-the-row spraying, a guidance system to maintain the implement over the pyrethrum rows, and a precision weed sensing system incorporating image analysis algorithms and machine vision equipment to recognize weeds from images and make the decision to actuate solenoids to spray weeds.

Methodology

This HIA project developed a precision weed sensing system to expand the pyrethrum industry's weed control methods and overcome limitations with existing spot spray technologies that are unable to differentiate between plant species that are growing side by side in the crop row. The methodology reported here is a summary of the project's Milestones 102 to 108 reports, which are included in Appendix A.

As per Milestone 102, BRA and NCEA determined design considerations and the project KPIs for the precision weed sensing system in an initialisation meeting held in 2012 after an initial assessment of a number of pyrethrum growing sites in Tasmania. A test data collection device was designed, built, tested and deployed on BRA pyrethrum sites. Video data was collected and reviewed by researchers at the NCEA and it was determined that the crop and weed features in the images were separable and suitable for classification (Milestones 102 and 103). In consultation with collaborators, the period between March and July was identified as suitable for undertaking spot-spray operations to address the industries weed control concerns. Dieback of the plants was identified as a limitation in the final two months of the spraying season.

The weed management approach undertaken during the development of the precision weed sensing system was two wide row bed-layouts that involved speedlings planted on 60cm rows and planting seed on four rows spaced at 45cm over the 1.8m wide bed (Figure 1.1). The wide row layout aimed to assist weed management by providing space for mechanical or chemical weed control methods to be applied between the rows and chemical control on the rows.

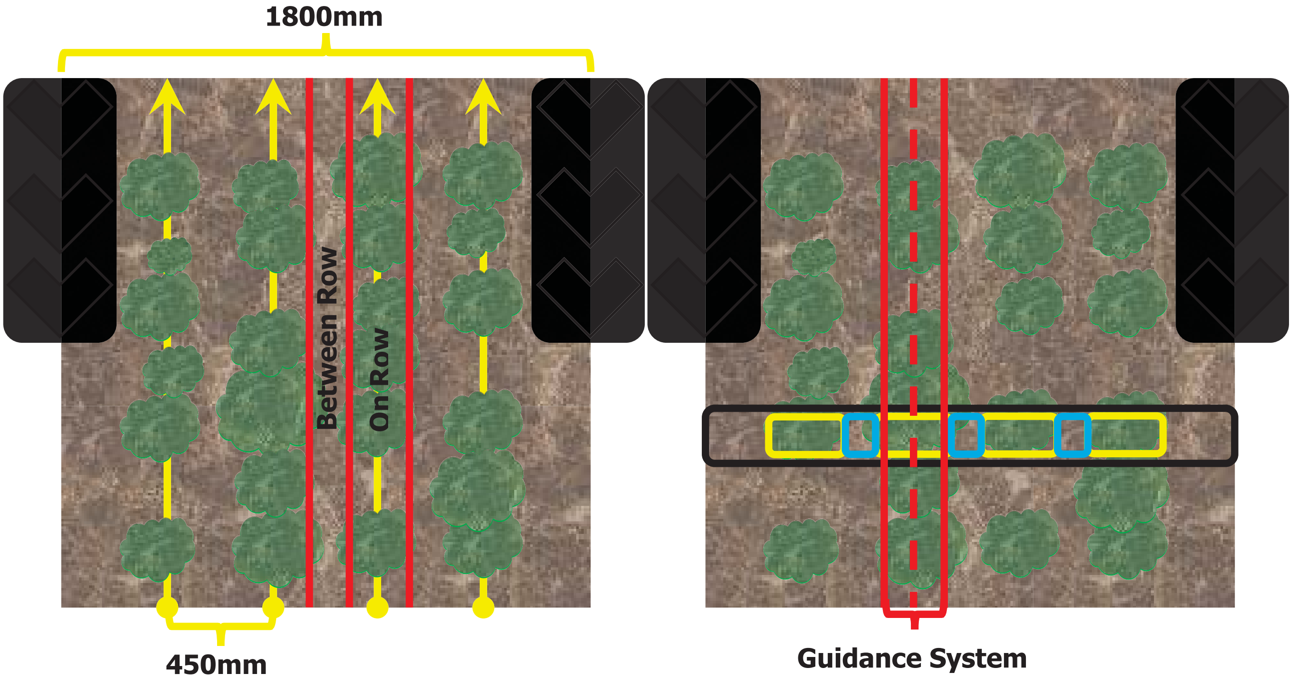

Figure 1.1: Example of the systems configuration in respect to the pyrethrum rows. Blue rectangles represent between-the-row regions of interest for spot spraying and yellow rectangles represent on-the-row regions for spot spraying

The wide row layout required the image analysis to perform three functions. In Milestone 104, a row detection algorithm was required to detect the row in the image so that on-row and between-row regions of interest in the image could be defined. A weed detection algorithm analysed these image regions to control on-the-row and between-row spray delivery systems separately. A row detection algorithm facilitated control of the implement guidance system to maintain the on-the-row spray delivery system over the row.

A prototype implement was assembled and evaluated to determine how effective it was at spot spraying weeds (Milestones 104 and 105). The prototype implement incorporated a precision spray delivery system for on-the-row spraying, spray shields for between-the-row spraying, a guidance system to maintain the implement over the pyrethrum rows, and a precision weed sensing system incorporating image analysis algorithms and machine vision equipment to recognize weeds from images and make the decision to actuate solenoids to spray the weeds.

Proof-of-concept spot-spray prototype-implement

The 2.6x1.8x1.5m spot-spray prototype-implement (Figure 1.2) consists of two detect-spray systems, a control computer and a control box to allow simultaneous weed detection and spraying of a bed of four rows (Milestone 104, 105, 107 and 108). The detect-spray systems each monitor two adjacent crop rows and each consists of:

- one camera which views two adjacent crop rows;
- two pairs of spray solenoids, with each solenoid covering half of one crop row; and
- one multi-core computer to analyse the images from the camera, trigger the spray solenoids and communicate with the control computer in the tractor cab.

A control computer with user interface is positioned in the tractor cab, and communicates with each detect-spray system via a CAN communications link. The control computer allows the user to configure software parameters of the detect-spray systems.

a) **b)**

Figure 1.2: a) The prototype attached to the back of the tractor b) A 3D CAD drawing of the implement with inter-row methods trialed attached to the toolbar.

On-the-row spraying

Two regions of interest (ROIs), over the two centre rows of the bed, were configured for real-time field evaluation of the 1.8 m multi-row prototype (Figure 1.3). The system was configured for the center two rows as variability in pyrethrum appearance on the outside rows of pyrethrum was increased by the tractor tyres damaging the pyrethum that grew into the wheel track. When a weed was detected within a region the spot-spray solenoids were instructed to open and spray Figure 1.4.

Figure 1.3: The rectangular regions of interest (ROI) specified over the rows of pyrethrum for on-the-row analysis

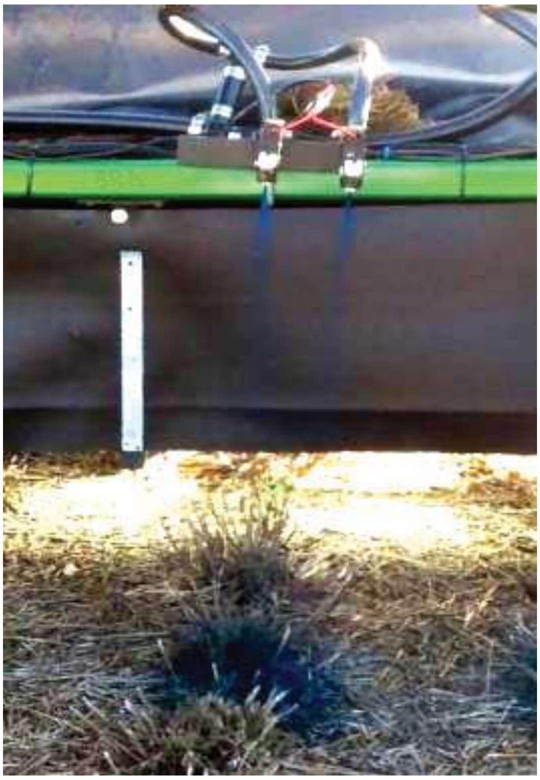

Figure 1.4: The spray system operating on-the-row

Between-the-row weed control

Two methods of weed control between-the-row were trialed (Milestone 105). The first method involved the design and manufacture of weed-control spray method with cutting discs to trim the pyrethrum in path of the inter-row shield sprayer and avoid spray hitting the pyrethrum and translocating to the remainder of the plant (Figure 1.5 (a)). The second weed control method consisted of a tyne and sweep to mechanically remove weeds between the row Figure 1.5 b). These inter-row weed control methods were trialed to determine how effective they were at removing weeds between the pyrethrum rows Figure 1.5 c).

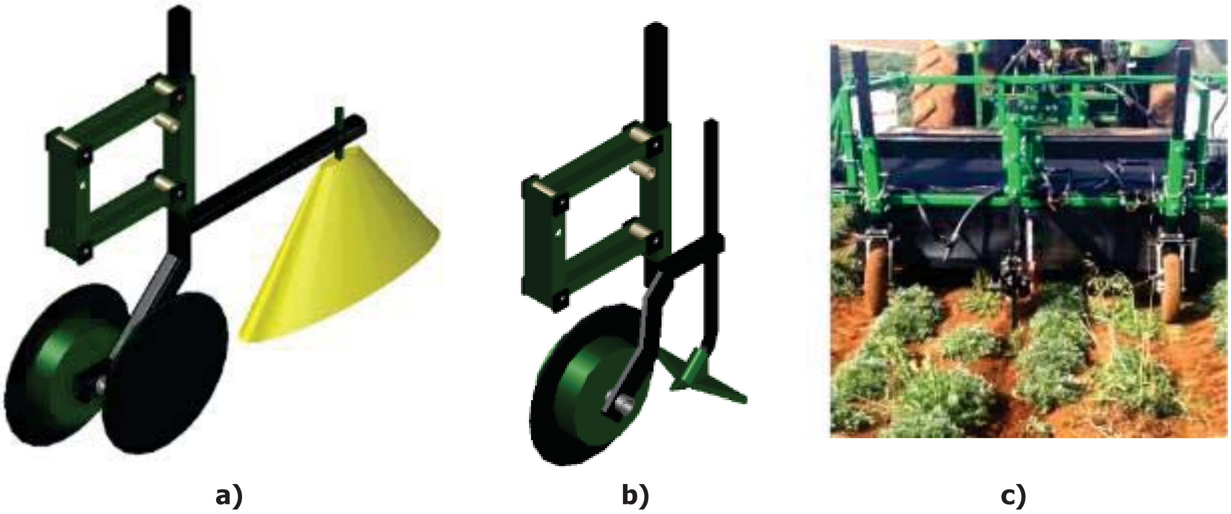
a) b) c)

Figure 1.5: a) Inter-row weed control spray method with cutting discs to avoid any spray drift translocating into the pyrethrum; b) Inter-row weed control mechanical method with tyne and sweep to remove weeds between the row; c) The mechanical weed control method in operation on the back of the implement

Guidance system

The wide-row planting required a guidance system to maintain the spray nozzles on-the-row and either tillage or shield-spray equipment between-the-rows. As the wide rows were generally well defined a vision guidance system was incorporated to maintain the implement in the correct position. An offset error from the row found by the row detection algorithm facilitated control of the implement guidance system (Milestone 105). When the offset deviated away from the centre of the pyrethrum in the image (Figure 1.6) the side-shift-hitch was instructed to move. Figure 1.7 shows the system detecting the row and isolating the between-the-row and on-the-row regions.

Figure 1.6: Image of the guidance line determined through image analysis to maintain the on-the-row spray system over the row.

Figure 1.7: Image of the guidance line (green) determined through two rows and bounds for on-row and between-row regions of the pyrethrum (yellow and white lines).

Precision weed sensing system

Image analysis algorithms were developed and evaluated that discriminated pyrethrum plants from weeds using colour and depth, and extracted features including colour, texture and shape of different plants (Milestones 103, 104, 107 and 108; extract of Steve Rees's PhD thesis 'Depth Colour Segmentation Algorithm', Appendix C). A processing technique was also developed that enables real-time operation at commercial groundspeeds (Milestone 103; extract of Steve Rees's PhD thesis 'Synchronised Parallel Processing', Appendix C). A demonstration video of pyrethrum classification is attached to this report submission.

As for the row detection algorithm, the weed detection algorithms had to cater for a wide range of conditions related to the weed orientation and position relative to the pyrethrum and the pyrethrum growing configuration (i.e. bed layout and planting). Refer to Appendix B for usage guidelines for the precision weed sensing that incorporates field conditions.

The weed spectrum included weeds that were sparse such as grasses, weeds that were high and dense (potatoes, mallow), weeds were low and dense (clover), weeds that had similar texture to pyrethrum (pink weed, bur-chervil Figure 1.6 (a)) and weeds with similar shape (pink weed).

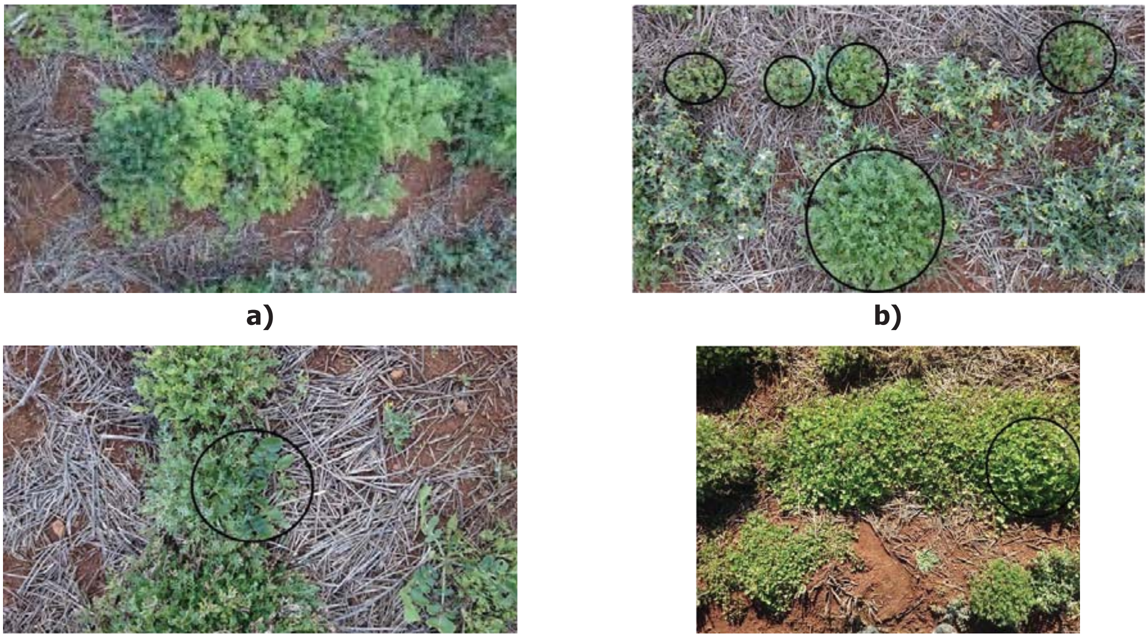

a) b) c) d)

Figure 1.6: a) Bur-chervil (lighter green) amongst pyrethrum plants canopy; b) Pyrethrum size ranging dramatically after regrowth (circled) in a groundsel infestation; c) Broadleaf weed butted up against a pyrethrum plant on-the-row; d) White clover infestation with an effected pyrethrum plant circled

Technology transfer, monitoring and evaluation, and project reach

The project has undertaken technology transfer activities, implemented a monitoring and evaluation framework, and had had local, regional, national and international reach as described below.

- The project has been conducted in close collaboration with BRA, from project development, to constructing prototypes and conducting field trials and evaluation (see Evaluation and Discussion section).
- The project has been presented at national and international scientific forums, national industry forums, local and regional field days for audiences consisting of the growers, consultants, chemical companies, agricultural machinery companies, policy makers and the general community (see 'Communications' in Outputs section below).
- Project outputs are part of a USQ commercialisation strategy for weed spray technology, which has also involved discussions with John Deere (see IP and Commercialisation section).

Outputs

The outputs as per the project plan are as follows, and each output is described further below.

1. Single row test rig and proof of concept algorithms
2. Field prototype, performance specification and usage guidelines
3. Variation: Pre-production spot spray unit
4. Final report
5. Communications

1. Single row test rig and proof of concept algorithms.

A single row test unit (Figure 2.1) was developed, manufactured and deployed with BRA for data collection. The unit and data collected is detailed in Milestone 102 and 103 (Appendix A).

Figure 2.1 Single row test rig in a field of pyrethrum.

The following machine vision algorithms were developed to identify pyrethrum from weeds for 60 cm rows and 45 cm rows:

- Spatial position of weed in row. This algorithm is novel but not patentable. The algorithm identifies weed plant from a crop plant by comparing the position of the centre of the plant under scrutiny against the centre of the row. If the plant under scrutiny does not have a similar centre position, it is a weed.
- Size and height. This algorithm is a software implementation of known machine vision concepts. The algorithm determines weeds plants that are larger or smaller or taller or lower than the bounds of a crop plant.
- Texture. This was a software implementation of an existing technique called Local Binary Patterns (LBP)
- **Patentable IP.** Two pieces of patentable technology (International Application PCT/IB15/1604) were developed in this project being:
    - A real-time Synchronised Parallel Processing (SPP) technique. This technique is an enabling technique to allow the parallel micro-processors to compute complex algorithms requiring significant execution times but still function in real-time.

- Depth Colour Segmentation Algorithm (DCSA). The DCSA is a machine vision segmentation algorithm that combines colour and depth data to identify individual leaves in an image even when occluded.

2. Field prototype, performance specification and usage guidelines.

In a project variation, delivery of the field prototype was replaced with delivery of a pre-production spot spray prototype. Performance specifications and testing of the algorithms (Table 1) was completed and is documented in Appendix milestone 103. The usage guidelines for use is in Appendix-Guidelines.

Table 1 Summary of the performance of different algorithms.

In-control weeds			
Feature extraction method	**Miss rate**	**False trigger rate**	**Hit rate**
DCS	1%	3%	93%
DCSS	1%	1%	98%
LBPD	1%	2%	92%

3. Variation: Pre-production spot spray prototype.

A pre-production spot spray prototype was developed, built and deployed in Tasmania (Figure 2.2). Specific outputs were researched and developed to allow the pre-production unit to be able to effectively spray weeds in the pyrethrum rows with minimal overspray. The specific outputs are described below.

Figure 2.2 Pre-production spray rig in a field of pre-harvest seed planted pyrethrum

Row guidance.

A colour row guidance algorithm, with plus or minus 5 cm accuracy, was developed in the project. The row guidance algorithm coupled with a side shift hitch maintained the pre-production prototype position accurately over the row. A side shift hitch is a 3 point linkage adapter plate that fits between the tractor and the sprayer (or implement) and allows the sprayer to move left or right in relation to the tractor. Correct alignment of the pre-production prototype, cameras and spray apparatus on the row is vital for the spatial position algorithm to function correctly and the spray application to be accurate.

Spot spray application apparatus domain knowledge.

Domain knowledge of spray application apparatus was gained from trials. This included the solenoid valve operation power consumption, on/off timing and delays in actuation as well as the spray nozzle spray pattern. The trials determined the effect different solenoid and nozzles had on the spray application to the weeds due to groundspeed differences.

Pre-production field ready prototype.

A pre-production prototype spot spray system was designed and built specifically for the above mentioned algorithms and spray technologies. The pre-production prototype sprayer included a light diffusing hood with a lighting system under the hood. The hood and lighting system excluded external ambient light and provided the camera with a constant illumination and colour light source.

System software.

A software program was developed to provide a graphical user interface to the operator and provide adjustments to the machine vision and sprayer system operation.

Plant mapping.

The system software saved the images and GPS co-ordinates as the system was operated. The saved data provides the information necessary to create plant maps and additional benefits outlined as data mining in the outcomes section.

Communications

Academic presentations

- Rees, S and McCarthy, C (2014) Occlusion-tolerant image segmentation using colour and depth data with application to weed spot spraying in sugarcane and pyrethrum crops, In: 5th International Workshop: Applications of Computer Image Analysis and Spectroscopy in Agriculture, 12-13 July 2014, Montreal, Canada.
  (Attached)

- Rees, S and McCarthy, C (2014) Real-time processing technique for image analysis-based weed spot spraying at commercial ground speeds, In: 5th International Workshop: Applications of Computer Image Analysis and Spectroscopy in Agriculture, 12-13 July 2014, Montreal, Canada.
  (Attached)

- USQ PhD thesis "Precision weed detection via colour and depth data fusion in real-time for automatic spot spraying" by Steven Rees
  (Excerpts attached, Appendix C)

Industry articles

NCEA's weed sensing research for pyrethrum funded by HAL has received media exposure in industry magazines as a result of our parallel research for the RIRDC and SRDC.

- The Land 'New weed tech just years away' (7 March 2013, p.40)
- Farmer's Weekly digizine (South Africa) 'Weeding out the bad guys' (December 2012)
- The weed spot spraying research funded by HAL featured prominently in the award application.
- Abstract accepted for paper presentation: Rees, S and McCarthy, C (2013) Development of machine vision-based precision weed sensing for interstate sites using broadband internet services, In: Digital Rural Futures Conference, 26-28 June 2013, Armidale, Australia.
- Rural Weekly Times newspaper in Southern Queensland.
- International Pest Control Magazine in the U.K.
- The Grain Grower Magazine.

Presentations

- LandWISE conference (New Zealand, May 2015)
- Burnett Mary Regional Group / Growcom Carbon Farming Project grower tour in Marburg (April 2015)
- USQ Professional Lecture for Engineers Australia, Toowoomba Regional Group (March 2015)
- Upper North Farming Systems Group Precision Agriculture Day (via Skype, February 2015)
- Lecture at IEEE Summer School on Agricultural Robotics (Sydney, February 2015)
- IEEE Agricultural Robotics and Automation (AgRA) Technical Committee webinar lightning round (January 2015)
- The research was represented at the NCEA stand at Heritage Ag Show in Toowoomba in September 2014.
- McCarthy, C (2014) Real-time machine vision applications for crop sensing, In: Proceedings of the 17th Symposium on Precision Agriculture in Australasia, 2-3 September 2014, Adelaide, pp. 86-87.

Radio interviews

A radio interview with ABC radio for northern Tasmania has been conducted.

Demonstrations

Demonstrations of the weed spot spraying equipment for pyrethrum have been given to the following visitors to the NCEA:

- Dr Geoff Garrett (Queensland Chief Scientist)
- Hon Ian Walker (State Minister for Science, Information Technology, Innovation and the arts)
- Hon Ian Walker's Chief of Staff
- Members of the Trade Investment Queensland Scholarship Familiarisation Agency
- High school students as part of USQ Equity Project
- Australian and international researchers visiting the NCEA

Farm days

Preliminary investigations have been presented at a BRA research day in Tasmania in July 2012.

In March 2015, Dr Matthew Tscharke presented the pyrethrum weeds work and future opportunities using the technology at the Northern Tasmania HIA ‘Between the Rows’ regional roadshow in Latrobe Tasmania.

Commercial partners

Dr Steven Rees, Dr Matthew Tscharke, Prof Craig Baillie and Dr Erin Rayment (USQ Commercialisation) undertook a discussion and presentation of the protected IP to John Deere senior management for the Intelligent Solutions Group (ISG) in Des Moines in July 2014.

Awards

Dr Cheryl McCarthy was a finalist in Queensland’s Young Tall Poppy Science Award in August 2015. The HIAL weed spot spraying research featured prominently in the award application.

Outcomes

The primary outcome of the project was to deliver performance specification and usage guidelines based on the proof-of-concept technology for automated spot spray weed control strategies. The performance specification and guidelines have been discussed in the previous ‘Outputs’ section. The specific KPIs required to meet the outcome are discussed below. Outcomes of reduced spray application and uses in weed mapping and new precision farming techniques as per the project plan are also discussed below.

Achievement of specific KPIs.

Intended outcomes.

The following KPIs were determined in consultation with BRA to indicate successful achievement of a proof of concept automated spot sprayer for identifying weeds in 60 cm and 45 cm pyrethrum rows:

1. Distinguish weeds at an operating speed of 6-8 km/h
2. Day (with shields) and night time operation
3. Typical window of operation from post-harvest to end of July
4. Detect weeds 100mm diameter with a 98% accuracy

Achievement.

1. The ability to distinguish weeds at 6 to 8 km/h was achieved with the enabling Synchronised Parallel Processing technique. The technique allowed complex, computationally intensive algorithms to be developed to identify the weeds in pyrethrum and operate in real-time at groundspeeds up to 8 km/h. Simulations of the algorithm showed that the machine vision system could operate up to groundspeeds of 17.2 km/h without missing areas of the ground but in practice this is not achievable as the physical ground conditions and machinery cannot operate at these groundspeeds.

2. Day and night operation was achieved by the addition of a light diffusing hood and a constant light source. The light diffusing hood maintained a controlled area where direct and ambient light cannot effect the image analysis and the constant light source supplied a known intensity and colour to enable consistent image analysis.

3. Window of operation. The window of operation was modified from March-July, to March-mid-May. This decision was made in consultation with BRA after ‘die-back’ was noticed in the collected data. Die-back is where part of the pyrethrum plant dies off, particularly in wetter years. The effect of die-back on the image analysis system was that the pyrethrum plant no longer was consistent in colour which changed the perceived shape and texture of the plant in the image and caused the system to identify it as a weed. By spraying before mid-May, the majority of die-back effects were missed as die-back became more noticeable in the colder months.

4. Detect weeds 100mm diameter with 98% accuracy. The approach taken in the algorithm development was to identify pyrethrum and discriminate all other plants as weeds. The algorithms developed after year one were able to identify isolated weeds that were 100 mm or less at 98% and better as shown in Figure 3.1. An issue encountered with the KPI was that pyrethrum and weed plants were not always isolated from each other and therefore other

algorithms were required to be developed. In 60 cm rows of pyrethrum a Depth, Colour, Size and Spatial (DCSS) algorithm was developed which achieved the required 98% accuracy for pyrethrum identification with a 1.2% overspray rate. Notably pyrethrum grows differently under various row spacings and features that were prominent in one row spacing, such as minimum pyrethrum plant size, is not as prominent in another. In comparison results for the 45 cm row spacing were 88% accurate with a 12% false positive rate. Therefore these results indicate that a 60 cm row spacing cropping system would provide the optimum spot spray result.

Figure 3.1. Weeds spot sprayed with blue dye by the developed system.

Project variation

The variation of the project added the development and deployment of a pre-production prototype. This was achieved as seen in Figure 1.2 with the prototype operating on 45 cm rows.

Additional outcomes benefits

1. The project has led to the development of the Synchronised Parallel Processing (SPP) technique and Depth Colour Segmentation Algorithm (DCSA). The SPP technique is a real-time enabler for machine vision in other areas such as sugarcane, cotton and grains. The DCSA is a versatile algorithm for addressing occlusion in images in other crops as well. HIAL has entered into a multiparty agreement with USQ and SRA to allow use of this IP for further development and commercialisation.

2. The saving of the colour and depth video data referenced to a latitude and longitude allows ‘data mining’ to be undertaken post process. Data mining is where the data can be ‘mined’ for information other that what it was recorded for. In this case plant health and plant density maps could be generated and evaluated.

Longer term outcomes

Longer term outcomes may be significant depending upon the results from data mining the video streams saved at the time of operation. This approach of post processing the information will fulfil some additional weed control techniques identified in the initial project document being:

- Weed scouting. The system might be used to automate manual infield weed scouting tasks.

- Generation of weed maps. Across several seasons, automatically-generated weed maps might be used to identify emerging weed problems / persistence, and thus inform crop rotation decisions and other weed management strategies.

Narrow row

An additional longer term outcome is the further development of the pre-production unit into 25 cm pyrethrum rows being undertaken independently between the USQ and BRA. At the start of the project, it was felt that the pyrethrum farming system could be modified to wide rows (60 cm and 45 cm) which would provide the capability to spray the weeds with limited overspray onto the pyrethrum. However over time this modification to the cropping system has not proven successful with operations reverting back to a 25 cm row spacing.

Detail all economic, social and environmental impacts

All trials in this project were achieved by spraying blue dye onto the weeds and ground as seen previously in Figure 2.1. From visual inspection of the blue dye applications, areas of the field reduced the spray used by 80% compared to a broadcast spray. There was negligible reduction in out of control weed situations, but this was less than 1% of the fields trialed. This is similar to the fallow spraying in the green from brown products (Weedseeker and Weed-it) where farmers found a savings of between 40 and 90% of herbicide used.

The results suggest that the use of the spot spray system (once commercialised) will benefit the environment, animals and farm workers by using less herbicide and reducing the amount of run-off and effects on animals, fish or farmers. The technology potentially enables use of alternative herbicides which will relieve the traditional herbicides and reduce the potential of herbicide resistance.

Evaluation and Discussion

Effectiveness of project activities.

The project activities were very effective in delivering the required outputs outlined above. The project was a collaboration between Botanical Resources Australia (BRA) and the University of Southern Qld (USQ) and centered in Ulverstone Tasmania. The distance from Toowoomba Qld (USQ) to Ulverstone did not pose any operating issues for continuous data collection over the post-harvest vegetative growth stage, as USQ used remote access to computer systems whilst the BRA agronomist undertook the field data collection. This was achieved by using a data collection device designed specifically for this application.

Algorithms were able to be developed in Toowoomba and the BRA agronomist was able to collect data to specifically test the algorithm. This sped up the algorithm development cycle at a reduced cost to the project and was only achievable due to the close working relationship between USQ and BRA. The BRA head of research (Sarah Pethyridge) left BRA during this project, but this did not affect the outcome as it was a team environment at both BRA and USQ on this project and leadership from BRA was undertaken by Tim Groome.

Unforeseen variations to planting configurations in multiple years of the project led to additional image analysis algorithm development and refinement. It was found that pyrethrum grew differently in different configuration, with image analysis features prominent in one situation not being prominent in another.

The algorithms were initially developed on row spacing of 60 cm, as per initial project meetings with BRA and anticipated benefits for weed control efficacy. In year two BRA undertook row spacing trials of 3 rows per 1.8 m bed and 4 rows per bed, with 4 rows (45 cm) being identified as the spacing to take forward for economic and agronomic reasons. The algorithm development was then changed to pre-harvest planted pyrethrum on 45 cm rows. The following year there was post-harvest 45 cm pyrethrum to trial on and the algorithm development was changed to 45 cm post-harvest. In the final year, BRA made the decision to remain on the narrow row spacing (25 cm) for a number of practical economic and agronomic reasons. BRA is now working with USQ separately to this project, to convert the algorithms over to the 25 cm row spacings.

Feedback

As mentioned previously there was a close working relationship between the BRA staff on the project and the USQ staff on the project. Due to this close working relationship practical feedback and field issues were communicated directly to the USQ team on a weekly basis during the post-harvest growth stage data collection and trial evaluation periods. The USQ team was then able to incorporate this feedback into the development of the algorithms, overall machine vision systems and pre-production prototype sprayer.

Additionally, USQ staff attended BRA/farmer days annually to explain to farmers what was being developed and the farmers were able to provide feedback to the researchers first hand. The feedback from the farmers mirrored that from BRA which would be expected as BRA is also the largest pyrethrum grower in Australia.

Feedback was also obtained at a stop/go meeting after a project variation was applied for to add the development of a pre-production prototype to the project. The stop go report is included in the Appendix.

Changes resulting from the project

The outcome of the project was a field prototype spot spray system for wide row pyrethrum (with 60 cm rows identifying as the optimum). This product has not been commercialised to date however the patentable IP from the project include enabling technology in the field of agricultural machine vision applications. The SPP technique will allow the spot spray and vision guidance systems to operate at common groundspeeds for the grains, cotton and sugar industries. This was not achievable using current processing techniques as they were too slow and limited the groundspeed of the systems.

The DCSA has a similar enabling effect in crops that exhibit high levels of leaf occlusion, which are generally the grass like crops. This will be of major benefit to the sugar, sorghum and corn cropping systems as it can identify the occluded weeds in the crop where presently no algorithm exists for this.

Outlined above, it was expected that the industry would change to wide row cropping systems, but this has not occurred. However, BRA is continuing to work with USQ to convert the wide row system over to narrow 25 cm rows and this will have significant impact on the pyrethrum industry when completed. The ability to spray weeds in narrow row pyrethrum will ease pressure off the selective herbicides used by the industry and allow operators to target hard to kill weeds. Once the conversion is completed, the same system could then be used to spray self-sown potatoes out of other vegetable crops.

Learning and relevance to industry

The learning from this project include:

- Spray application domain knowledge.

  The testing of existing spray apparatus associated with spot spraying identified areas of poor performance specifically on nozzle tips with large fan angles. The larger the fan angle, the greater the distance before and after the weed that the nozzle was required to be on. This led to using several 15 degree nozzles rather than one 40 degree nozzle. This is an issue in spraying weeds in crop as the larger spray footprint will damage the crop more.

- Usage guidelines for the system in pyrethrum including when to use it.

  The project found that pyrethrum grows differently in differing row configurations. This means that an algorithm will not necessarily work the same at all growth stages and therefore, there are times through the season that the system will be more effective than others. The guidelines outline what to look for in the regrowth to identify the optimum time to operate the spot spray system.

- How to operate a data capture intensive project where large distances are involved.

  A significant learning was how to develop systems and software to cater for remote access of researchers as well as how to co-ordinate with onsite agronomists and staff. This was achieved by the use of digital technology. The effectiveness of real-time interaction was reduced in areas of poor connectivity but not stopped.

Recommendations

Research

Follow on research should be targeted in 4 areas:

1. Application of the precision weed system to the wider horticulture industry. The developed system could have a wider use in vegetable crops as well as orchards.
2. Development of a market entry algorithm and a generic sensor module. This is being progressed in a current SRA project. There is potential for HIAL to fund evaluations and refinement for pyrethrum, and other horticultural crops.
3. Narrow row use in pyrethrum. This is already underway with work being separately undertaken between USQ and BRA.
4. Data mining. Research into recorded data in pyrethrum and also vegetables may be able to post produce weed, crop health, crop density and biomass maps for use in a precision agricultural system.

Exploitation of IP

Commercialisation of the IP generated in this project is being undertaken by USQ in conjunction with SRA and HIAL.

Dissemination

After a commercial partner has developed an off shelf sensor, evaluations can be undertaken in the horticultural industry to determine what individual industries would benefit from its use.

Scientific Refereed Publications

None to report. Academic and industry presentations for the project are listed earlier in the Outputs section under ‘Communications’, including Dr Steven Rees's PhD thesis.

Intellectual Property/Commercialisation

USQ has developed a commercialisation strategy that proposes to incubate the technology with John Deere. The commercialisation strategy is based on advance discussions with John Deere's ISG group and the world wide group leader of incubation within Deere. In parallel discussions have also been held with a number of RDCs to bundle intellectual property from various projects and to progress via incubation under a multiparty agreement. Comments from SRA, HIAL and CRDC were received and have been incorporated into the multiparty agreement. A program of work has commenced in the sugar industry and funded by Sugar Research Australia. HIAL has also signed this agreement due to the contribution of Background IP (developed in horticulture).

The project has generated the following commercial IP:

- A real-time Synchronised Parallel Processing technique. This technique is an enabling technique to allow the processors to compute complex algorithms requiring significant execution times but still function in real-time. International Application PCT/IB15/1604.

- Depth Colour Segmentation Algorithm (DCSA). The DCSA is a machine vision segmentation algorithm that combines colour and depth data to identify individual leaves in an image even when occluded. International Application PCT/IB15/1604.

- Spatial position of weed in row. This algorithm is novel but not patentable and determines if a plant is a weed by the position of the centre of the plant against the centre of the row.

- Size and height. This algorithm was a software implementation of known machine vision concepts. The algorithm determines weeds that are larger or smaller than pyrethrum and taller or lower that pyrethrum.

- Texture identification of pyrethrum. This was a software implementation of an existing technique called Local Binary Patterns (LBP).

Acknowledgements

BRA has been a close collaborative partner in this project. Specifically the efforts of Francis Chamley, Tim Groome, Simon Wilson, Lynden Head, Sarah Pethybridge and the maintenance harvest crew staff have been integral in achieving the outcomes in this project.

Appendix A – Detailed report for each milestone

Milestone 102

The project is progressing satisfactorily with all objectives in Milestone 101 being met. A project initialisation meeting was held at Botanical Resources Australia (BRA) headquarters in Tasmania along with a site visit and assessment of BRA pyrethrum paddocks. The site assessment allowed the BRA and NCEA team to determine the KPIs for the weed detection system as well as identify any design considerations that should be addressed. A single row test data collection device was designed, built, tested and deployed at BRA with BRA staff collecting data for analysis. The data provided has enabled the NCEA to determine that the images should have enough separable features for classification and a prototype algorithm has been developed. Dieback of the plants has been identified as a limitation but this should only effect the final two months of the spraying season. The spraying season for the spot spray system will now be from March to July which is considered adequate to address the weed control concerns of the industry.

Develop single row test rig for preliminary data collection

A single row test rig has been developed which is capable of collecting data with a range of machine vision sensors. The test rig design has used findings from an earlier NCEA machine vision project for weed identification, which found that occlusion and illumination are major sources of error for machine vision and that the addition of a light diffusing hood and constant illumination can significantly reduce the errors. The lighting chosen is determined by the sensor being used, e.g. the colour and depth camera uses white LED light so that there is no near infrared (NIR) light being emitted that can interfere with the depth sensing. The test rig was designed to fit in the back of a utility and be pushed down the rows. Figure 1.1 (a) shows the test rig in a field of pyrethrum (Note the enclosed hood and the transportation mode (push)). Figure 1 (b) shows one example of the lighting and sensor systems available. This sensor mounts on top of the rig in Figure 1.1 (a).


(a) **(b)**

Figure 1.1: The test rig (a) in a row of pyrethrum and (b) the LED lighting assembly with Kinect colour/3D sensor

Data collection of crop and weed under a range of agronomic, meteorological and operational conditions

An image acquisition system was developed which consists of a fitPC (a low power, low form-factor computer), computer monitor, mouse, keyboard and sensor system. All the computer equipment is solid state and robust to field use. A software application was developed to acquire the colour and depth data and store this data to the computer hard drive. This system was capable of a playback system for verification of data.

Data was collected by BRA from April to August on approximately a three week interval and transferred to the NCEA via Dropbox which is a site sharing service for data (e.g. files and photos). Data has been collected for six sites which comprise a comprehensive range of conditions and weeds in which the pyrethrum grows. The requirements used to guide selection of sites and schedule for data collection.

Site criteria for data collection

The sites should have the variety of weeds that will need to be discriminated against mixed in with the pyrethrum and the sites should be indicative of the main soil types and stubble cover. The ideal would be to have sites that cover the main soil types e.g. black, brown, red, sandy, with each soil type having sites that have the differing stubble cover, e.g. trash, no trash and each stubble cover having the weed spectrum in it.

Trial criteria for evaluation

Set up trial runs that cover the criteria and acquired footage once per week. If a collection time is missed because of rain or other weather condition than the soonest the footage can be acquired will suffice.

The purpose of the collection from the sites is to get footage of the weeds and pyrethrum plants at all growing stages in all conditions i.e. ground cover and soil type so that effective algorithms can be determined that work in all conditions or if they work at a specific growth stage that also can be determined.

Testing

The algorithms will need to be tested on the sites that cover the main soil types e.g. black, brown, red, sandy, with each soil type having sites that have the differing stubble cover e.g. trash, no trash, golden stubble and weathered stubble, and each stubble cover having the weed spectrum in it at different sizes.

If the algorithm is to run in the daylight then a test of its consistency (trigger size, shape, and colour) over varying daytime conditions needs to be achieved along with the variations outlined above. That is calibrated when a cloud is overhead and tested and then tested in the direct sunlight and the differences in operations noted.

In general the most difficult segmentation will occur on black soil with golden stubble therefore this is one criterion that should be tested.

Develop KPIs for evaluation

An initialisation meeting and site visit was held at BRA to start the project and to undertake an initial assessment of the sites. Considerations for design of the precision weed spot spray system that were identified as application of herbicide and weed discrimination.

Herbicide application considerations for BRA were:

Plant density versus row spacing. The physical capacity to precisely spray a weed that is growing next to and/or touching the pyrethrum plant without affecting the pyrethrum plant has not been developed at this time. A solution was to use a wider row spacing which would allow application technology used in the row crop industries to be used.

Machinery aspects: This includes larger multiple hood design required to allow daytime use of a camera system outlined above in the test rig design. The tractor and spray rig is required to be precise within the rows of pyrethrum without affecting the pyrethrum.

Herbicide interactions in the row interspace with the crop. Look to trial two rates, i.e. a heavy application rate might be acceptable for spot spraying larger weeds and a lower application rate blanket sprayed for small and emerging weeds.

Herbicides identified that can be used on the crop row. Herbicides will be trialed to determine the damage to the crop from overspray and drift

NCEA will need to consider the following KPIs for weed discrimination:

- Distinguish weeds at an operating speed of 6-8 km/h
- Day (with shields) and night time operation
- Typical window of operation from post-harvest to end of July
- Detect weeds 100mm diameter with a 98% accuracy

Weed detection will most likely include two algorithms by partitioning the problem into weed detection strategies for:

- Over the row algorithm (green from green)
- Over the inter row (green from brown)

Analyse data to identify discriminatory attributes between crop and weed species. Develop detection algorithm

The captured footage was evaluated to determine discrimination ability. From visual inspection (typical image Figure 1.2 (a) it could be seen that there were height features that could be used as well as texture differences in the colour image. In Figure 1.2 (a), the input colour image contains the pyrethrum plant in the centre and a weed located to its top right. Figure 1.2 (b) displays the image in which non-pyrethrum pixels have been identified (rendered in yellow) using a prototype algorithm.

(a) (b)

Figure 1.2: The analysis process with colour and depth (a) colour image of crop and weed and (b) identified pyrethrum pixels plotted in pale pink and other coloured pixels indicating weed areas.

Deterioration in the pyrethrum plant was captured in the footage over time. This ‘dieback’ can be seen by comparing the greenness between April and July shown in Figures 1.3 (a) and (b) respectively. The dieback is being investigated by BRA staff and the short term solution for the project is to undertake the majority of our field work before July.

(a) (b)

Figure 1.3: Pyrethrum (a) in April with consistent green colour and (b) in late July with grey areas showing plant 'dieback'.

Milestone 103

The project is progressing satisfactorily with the collection of field data now complete. Further development of image analysis algorithms has been undertaken and has demonstrated that existing image analysis techniques are not satisfactory for discriminating weed from pyrethrum. Hence, custom algorithms for weed and pyrethrum discrimination have been developed and refined with promising results. A groundspeed of 8 km/h has been identified as the KPI for minimum groundspeed and we have developed a new processing technique which will allow groundspeeds of 8 km/h and above. The new processing technique and algorithms are currently being evaluated for IP protection and commercialisation. Initial evaluation of shield spraying between the rows has been undertaken and the results very positive. The results from the algorithms and shield spraying have highlighted the need for a more precise dosing system for the spot spray and an accurate guidance system for the shield sprayer as well as drift control. We have submitted a project variation request to HAL which proposes to conduct this additional research as part of the current project.

Evaluation of standard spot spray solenoid and nozzle assemblies have indicated that off-the-shelf products are suitable for spot spraying isolated weeds, but would cause overspray of pyrethrum if the weeds are growing close to the crop. A significant proportion of weeds grow amongst the pyrethrum plants so there is a need to develop and evaluate a more precise dosing system.

An initial evaluation for shield spraying of weeds between the rows was undertaken with very positive results. The results indicate that shield spraying could potentially provide effective between-row weed control in the pyrethrum industry. However, the trials highlighted the need for a satisfactory guidance system for the shield sprayer as well as the highest possible drift control. This additional research has been submitted to HAL as a project variation request for the current project.

Data collection of crop and weed under a range of agronomic, meteorological and operational conditions

Field data has been collected by BRA staff for a range of pyrethrum growing conditions and weed infestations. Data has also been collected by USQ staff under field operating conditions at expected working speed (up to 8 km/h) and at 25 frames per second, to enable evaluation of the algorithm's real time performance. The data was collected by USQ staff during a field visit to BRA in May 2013.

Video data has been collected under a range of growth stages and conditions in April-August 2012 and April-May 2013 for five sites in total (Tables 2.1-2.4). Crop plant diameter was in the range 15-30cm for all the data. The Greg Gibson site was a trial plot of 45cm rows.

Table 2.1: Weed species at each site, determined by scouting and/or inspection of data and field history.

Site	Unique weeds present at site	Weeds common between sites
BRA Jamisons	Hemlock	
DRF Speedlings	Wild radish, wild carrot / knotted hedge parsley, cleavers, prickly ox tongue, blackberry	Flat weed, groundsel, thistle, sow thistle, dandelion, white clover, wireweed, red clover
Coles	Wild radish, wild carrot, subterranean clover, fumitory	
Dick	Cleavers, potato, grass, scotch thistle	
Greg Gibson	To be determined by visual inspection of video	

Table 2.2: Data collected for growth stage PHVG (post harvest vegetative growth).

Date	Time	Location	Length of run (m)	Ground speed (km/h)	Crop height (m)	Plant density (m^2)	Sensor height (m)
12/04/2012	11:40-11:45	DRF Speedlings	40	3	0.10-0.15	15-20	1
26/04/2012	12:10-12:25	DRF Speedlings	40	3	0.15	15-20	0.85 & 1
11/04/2013	15:05-15:15	DRF Speedlings	60	6 & 8	0.15-0.20	15-20	1
11/04/2013	15:15-15:45	DRF Speedlings	280	6	0.15-0.20	15-20*	1
12/04/2012	10:15-10:20	Coles	15	3	0.10-0.15	5-10	1
12/04/2012	13:10-13:15	Dick	15	3	0.10-0.15	15-20	1
17/05/2013	15:15-15:25	Greg Gibson	140	3	0.15-0.20	15-20	1

* some parts with no pyrethrum

Table 2.3: Data collected for growth stage PHSD (post harvest semi dormant).

Date	Time	Location	Length of run (m)	Ground speed (km/h)	Crop height (m)	Plant density (m^2)	Sensor height (m)
24/05/2012	11:05-11:20	BRA Jamisons	15	3	0.20	15-20	0.85 & 1
24/05/2012	12:00-12:20	DRF Speedlings	40	3	0.20	15-20	0.85 & 1
5/06/2012	12:55-1:10	DRF Speedlings	40	3 & 7	0.20	15-20	0.85 & 1
5/06/2012	10:05-10:15	Coles	15	3	0.20	5-10	0.85 & 1
5/06/2012	13:40-13:50	Dick	15	3	0.20	15-20	0.85 & 1

Table 2.4: Data collected for growth stage PHD (post harvest dormant).

Date	Time	Location	Length of run (m)	Ground speed (km/h)	Crop height (m)	Plant density (m^2)	Sensor height (m)
26/06/2012	11:10-11:35	DRF Speedlings	120	3	0.20	15-20*	0.85 & 1
19/07/2012	11:25-11:45	DRF Speedlings	80	3	0.20	15-20	0.85 & 1
6/07/2012	14:25-14:35	DRF Speedlings	40	1, 2 & 3	0.20	15-20	1
9/08/2012	11:50-12:10	DRF Speedlings	80	3	0.20	15-20	0.85 & 1
18/05/2013	11:05-12:20	DRF Speedlings	280	3	0.20	15-20^	1
26/06/2012	10:00-10:10	Coles	30	3	0.20	5-10	0.85 & 1
19/07/2012	10:20-10:35	Coles	30	3	0.20	5-10^	0.85 & 1
9/08/2012	10:30-11:05	Coles	30	3	0.20	5-10	0.85 & 1
26/06/2012	11:55-12:05	Dick	30	3	0.20	15-20	0.85 & 1
19/07/2012	12:15-12:25	Dick	30	3	0.20	15-20	0.85 & 1
9/08/2012	12:25-12:35	Dick	30	3	0.20	15-20	0.85 & 1

* some parts are weed free, weeds only between the row, or weeds only within the row

^ some parts with low, sparse, or no pyrethrum

Further analysis of data to identify discriminatory attributes between crop and weed species.

The collected data has been used to develop image analysis algorithms for the purpose of discriminating between crop and weed. Development of image analysis algorithms comprises steps for segmentation, feature extraction and classification, as described below.

Segmentation: the analysis process of dividing the image into regions that are plant and non-plant. Ideally, if there is more than one plant in the image, the segmentation process will separate the plants for individual analysis. Depth and colour segmentation was found to effectively split the image into plants and a custom segmentation algorithm has been developed and further refined.

Feature extraction: the calculation of unique patterns and/or statistics for the different plant regions of the image, to be used as discriminatory features between the plant regions.

Classification: the categorisation of features extracted from the plant regions as crop or weed.

Texture classification

Visual inspection of the pyrethrum images indicated a textural difference between weeds and crop. An evaluation of texture analysis methods was conducted to determine whether texture was a satisfactory means of discrimination between crop and weeds. The evaluated methods were identified from a literature review and consisted of the Grey Level Co-occurrence Matrix (GLCM), Fast Fourier Transform (FFT), Grey Level Run Length Matrix (GLRLM) and Linear Binary Patterns (LBP). Each method returned different texture features. The techniques selected were rotationally invariant and illumination invariant. Illumination variance was minimised by enclosing the region of interest in a hood with a known constant light source.

Images were manually segmented into smaller sub-images of weed and crop for preliminary testing of the texture analysis methods. Twenty-five images each of crop and weed (all species) were used for a training set, and an additional 25 images each of crop and weed were used for a test set. The outputs of the texture analysis methods were stored in a csv file and further evaluated using Receiver Operating Characteristic (ROC) curve to identify the best features for discrimination. The best features were applied to classifiers for identification. The evaluated classifiers were a Support Vector Machine (SVM), Multi Layer Perceptrons (MLP), K Nearest Neighbour (KNN) and Standard Bayes.

Software implementation

A software program was written in C/C++ using OpenCV libraries to apply the texture analysis methods to an image and store the extracted feature results in a csv file. The software program also evaluated different window sizes (i.e. scales of texture patterns) ranging from 3 to 25 pixels. The texture analysis methods were applied as the window size was iteratively increased in steps of 2 pixels.

A Matlab script was written to read the results of the texture analyses and apply the t-test and ROC test. Execution of the Matlab script resulted in identification of the best three features from each texture analysis method and the best window size. The results from the Matlab script were then used as input for a classification application. The best three features from each texture analysis were used individually and then in combination to determine whether combining the texture analysis methods influenced or improved the classification results. The classification application was written C/C++ and OpenCV.

The field implementation of the image analysis algorithms is in C/C++ and OpenCV. Hence, use of C/C++ and OpenCV for evaluating the texture analysis methods and classifiers will enable the same results to be replicated in the field implementation, and minimise differences in results caused by using different languages, programming environment, compilers or libraries.

Texture classification results

ROC values for different combinations of texture features are presented in Table 2.5, and rates for false weed identification (False Positives), false weed misidentification (False Negatives) and correct weed identification (Hit) for different classifications based on texture features are presented in Table 2.6. The accuracies range from 44.1% to 61.8%. These accuracies are not satisfactory for an implementation into a real world system.

Table 2.5: ROC values for different combinations of texture features.

Test number	Texture analysis	Feature #1	ROC	Feature #2	ROC	Feature #3	ROC	Window size
1	GLRLM	LGRE	0.12	RP	0.07	HGRE	0.07	19
2	GLCM	CT	0.14	CS	0.12	MP	0.11	27
3	LBP	NU	0.25	NE	0.23	CE	0.14	15
4	GLRLM	LGRE	0.12	RP	0.07	-	-	19
	GLCM	CT	0.14	CS	0.12	-	-	27
5	GLRLM	LGRE	0.12	RP	0.07	-	-	19
	LBP	NU	0.25	NE	0.23	-	-	15
6	GLCM	CT	0.14	CS	0.12	-	-	27
	LBP	NU	0.25	NE	0.23	-	-	15
7	GLRLM	LGRE	0.12	RP	0.07	-	-	19
	GLCM	BT	0.14	CS	0.12	-	-	27
	LBP	NU	0.25	NE	0.23	-	-	15

NB. Features are:

CE = Centred Edge
NE = Number of Edges
NU = Non Uniform
RP = Run Percentage
CT = Cluster Tendency
CS = Cluster Shade
MP = Maximum Probability
LGRE = Low Grey-Level Run Emphasis
HGRE = High Grey-Level Run Emphasis

Table 2.6: Classification results (%) for different classifiers and combinations of texture features.

Test number	SVM			MLP			KNN			Bayes		
	FP	FN	Hit	FP	FN	Hit	FP	FN	Hit	FP	FN	Hit
1	7.9	34.2	57.9	34.2	15.8	50.0	13.2	23.7	63.2	13.2	28.9	57.9
2	17.6	35.3	47.1	0	50	50.0	17.6	35.3	47.1	5.9	44.1	50
3	26.3	18.4	55.3	21.1	21.1	57.9	31.6	15.8	52.6	23.7	23.7	52.6
4	35.3	17.6	47.1	44.1	11.8	44.1	20.6	26.5	52.9	0	38.2	61.8
5	21.1	34.2	44.7	23.7	21.1	55.3	23.7	21.1	55.3	21.1	23.7	55.3
6	23.5	26.5	50.0	29.4	20.6	50.0	23.5	26.5	50.0	8.8	29.4	61.8
7	26.5	23.5	50.0	29.4	23.5	47.1	23.5	23.5	52.9	5.9	32.4	61.8

Note: FP = False Positives, FN = False Negatives (i.e. misses)

Custom algorithm using colour and depth

A custom image analysis algorithm for weed identification has been developed that incorporates shape and size information from the collected colour and depth images, and typical results are presented in Figure 2.1 and the attached 10 second video. The image is split into three areas; left, centre and right. The left and right areas are analysed with a green from brown algorithm (i.e. kill anything green) with size selectivity, to allow the operator to set the minimum weed patch size to be targeted. The centre area runs the custom colour and depth algorithm. Pixels are highlighted in yellow when the pixels are identified as weed with a size in excess of that set by the user.

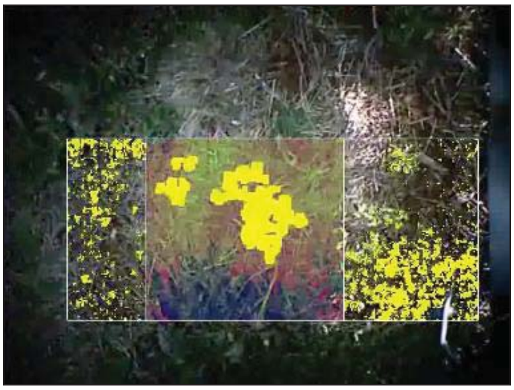

Figure 2.1: Typical results of the custom algorithm, with pixels rendered in yellow indicating weeds. Results are also provided as a video file (attached).

A finding from the current work is that shape is not a reliable discriminatory feature because the shape of the pyrethrum and the weeds vary and are not consistent. Texture was observed to be more consistent than shape. However, texture algorithms and classifiers from published literature were evaluated and found to be not satisfactory for discriminating pyrethrum and weeds.

Refinement of detection algorithm

Height sensing from depth images was evaluated and found to have some discrimination ability as the pyrethrum grew quickly to a medium size of approximately 20cm high and 30cm in diameter and then lay dormant for three months until spring. A custom analysis algorithm was developed which incorporates depth, spatial positioning (i.e. row geometry, since pyrethrum is planted in straight rows), plant size and range. The result from this combination of features is superior to our evaluations of existing published techniques for feature extraction.

The developed detection algorithm is required to operate in real time, which limits processing to a short period of time. Existing published techniques for crop and weed discrimination exceed the maximum allowable processing time and are not real time systems. Hence, we have developed a new multiprocessor technique which frees up the processing time ‘bottleneck’ of typical machine vision systems and will allow real time image analysis at groundspeeds in excess of 8 km/h.

Milestone 104

The project has been varied from the original to include the development of the spraying apparatus and between the row weed control system as well as the in row spot spray technology. This occurred at the end of Milestone 103 report and is reflected in the new achievement criteria for the project. Development has progressed satisfactorily according to the tasks set out for Milestone 104.

The field ready prototype has been designed and built to provide a platform for the evaluation in Milestone 105 as well as the larger field trials in Milestone 106. A number of different between row pieces of equipment have been designed and built to evaluate both herbicide and non-herbicide options. The spray delivery system has been developed and built to provide a highly optimised version of the commercially available technology. The row guidance system for keeping the prototype ‘on target’ has also been developed by using the images and results from the weed recognition software. This is expected to offer a low cost guidance solution for the pyrethrum industry and is ready to be integrated into the prototype in Tasmania. Further image analysis refinement is ongoing.

Development of the precision spray delivery system, spray shields for between row spraying and guidance system

Spray delivery system

Evaluations of commercially-available spray delivery systems (i.e. solenoid valves and nozzles) with a high-speed camera determined that there are time delays between spray triggering and establishment of a satisfactory spray pattern. The following optimisations were identified to reduce the time delay in the spray delivery system: a customised nozzle body to reduce residual pressure; and a trigger circuit for the solenoid valve to ‘over drive’ the turn on and quickly collapse the turn off power surge. The Asco solenoid valve exhibited the least time delay in the evaluations and was selected as the solenoid valve on which to implement the optimisations.

The aim of this research is to optimise the spray delivery technology that is currently available commercially on the market. In order to do this, a selection of solenoid operated spray nozzles used in agriculture for spot spraying were purchased along with a variety of spray nozzles. The different combinations of solenoids and nozzles were then evaluated by using a Photron SA3, high speed camera (frame rate of 2000 fps) to time the ‘turn on’ and ‘turn off’ delays of each system as well as the delay taken to create a required spray pattern. i.e an 80° flat fan nozzle will produce a spray pattern which has a flat fan shape and the fan has a width of 80°.

Time delay for satisfactory pattern establishment

A temporary enclosure was constructed to house the spray nozzles and contain the fluid from them. The enclosure consisted of a stainless steel water tray on the bottom with a drain hole in one corner. Three sides of the tray had Perspex mounted on them to a height of 80cm with the front of the area open so that there was nothing that could distort the images taken by the high speed camera. The solenoid nozzles were mounted in the enclosure 50 cm vertically above the centre of the floor. The camera was mounted approximately 1.5 m in front of the enclosure with a lighting module to provide lighting onto the spray nozzle (see Figure 3.1). The fluid used in the test was water that was dyed with blue food dye (40:1 water to dye). The blue food dye provided a contrast between the spray pattern and the background. The times that were measured for the evaluation were:

- The time delay between the fluid starting to leave the nozzle and the first droplet to land in the centre position on the floor (time centre).
- The time delay between the fluid starting to leave the nozzle and the first droplet landing on the edge of its spray width (time full band width).
- The time delay between the fluid starting to leave the nozzle and when the fan broadens to maximum spray width (time full band depth).

Table 3.1 and Table 3.2 below show the results of testing with a 40° nozzle and an 80° nozzle respectively. It can be seen from both tables (taken at 4 bar) that the time taken to reach the floor in the centre was similar in both nozzles at around 30 ms. The droplets reached the edge of their spray width quicker on the 40° nozzle (between 33 ms and 35 ms) than on the 80° nozzle (between 36 ms and 44 ms), which is expected as the droplets in the 40° nozzle have less distance to travel than the droplets in the 80° nozzle. Surprisingly, the 80° nozzle filled out to its full fan width more quickly than the 40° nozzle (between 50 ms and 58.5 ms, compared with between 55 ms and 64.5 ms). This is expected to be due to the larger orifice size in the 80° nozzle.

Table 3.1: Time measured for fluid to reach the floor from 50cm height with a Goyen solenoid spray valve and a 40° nozzle.

	Run One			Run Two			Run Three		
Pressure (bar)	Time Centre (ms)	Time Full Band Width (ms)	Time Full Band Depth (ms)	Time Centre (ms)	Time Full Band Width(ms)	Time Full Band Depth (ms)	Time Centre (ms)	Time Full Band Width (ms)	Time Full Band Depth (ms)
4.0	30	35	64.5	30.5	33	55	29.5	35	61
3.0	33.5	38	68	34.5	40.5	64.5	34	35.5	56

Table 3.2: Time measured for fluid to reach the floor from 50cm height with a Goyen solenoid spray valve and an 80° nozzle.

	Run One			Run Two			Run Three		
Pressure (bar)	Time Centre (ms)	Time Full Band Width (ms)	Time Full Band Depth (ms)	Time Centre (ms)	Time Full Band Width (ms)	Time Full Band Depth(ms)	Time Centre (ms)	Time Full Band Width (ms)	Time Full Band Depth (ms)
4.0	29	41	57	32.5	44	50	28.5	36	58.5
3.0	34	45	60.5	33.5	51	59	34	48.5	57

The reduction in time for the fluid to get to the maximum spray width from the 80° nozzle to the 40° led to the trialing of a 15° nozzle to evaluate if the time would continue to reduce . The 15° nozzle had a time of 30 ms to reach the floor as did both the 40 and 80° nozzles. However, as the fan was a very tight pattern the droplets reached the edge and also full depth at the same time as the droplets reached the floor. This indicates that a 15° nozzle can reduce the time delay for satisfactory spray pattern development from around 60 ms to 30 ms.

'Turn on' and 'turn off' times and residual pressure

A Light Emitting Diode (LED) was attached across the solenoid to indicate when the solenoid had been switched on and off. Application of power to the solenoid caused the LED to illuminate at the same time. The response time of the LED was 90 ns, compared with 0.5ms for each frame of the camera (which had a frame rate of 2000 frames per second). Hence, the delay times of the solenoid could be calculated to the nearest 0.5ms.

The turn on time was measured from when the LED was lit to when fluid started to emerge from the nozzle. Turn off time was from when the LED stopped being lit to when the fluid stopped emerging from the nozzle. When measuring the turn off times it was noticed that the solenoid valves fell into two categories which were:

A direct cut-off: This is where the fluid stopped emerging from the nozzle with a clear cessation.

A collapsed cut-off: This is where the fluid slowed to a dribble and stopped.

Figure 3.1 shows the LED (top, centre left of image) lit which indicates that the solenoid has had the power applied to it, but the fluid has not started to flow yet. This is the 'switch on' time. Figure 3.2 shows the spray pattern of an 80° nozzle immediately after the fluid has started to flow out the nozzle. The leading line of droplets can be seen half way down the image on the left.

Figure 3.3 shows the initial deactivation of a Goyen solenoid with 80° nozzle. It can be seen that the LED in the top centre left of the image is off and the width of the fan is not 80° and therefore has reduced in width and not directly cut-off. Figure 3.4 shows the dribble at the end of the deactivation period.

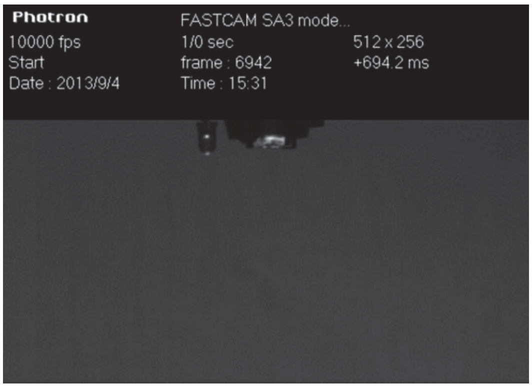

Figure 3.1: Turn on period when the solenoid is energised (LED is illuminated) and the fluid has not started to flow yet.

Figure 3.2: Spray pattern of an 80° nozzle after the fluid starts to flow after switch on.

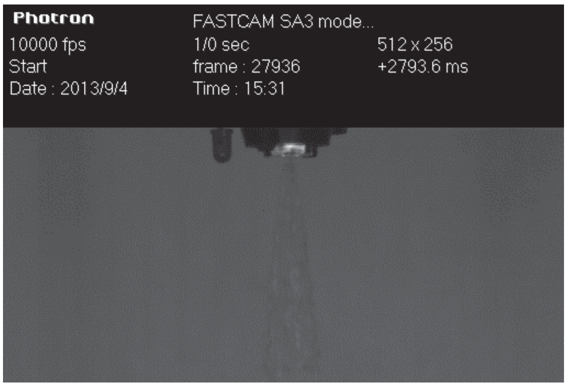

Figure 1.3: Initial deactivation of a solenoid valve exhibiting residual pressure build-up.

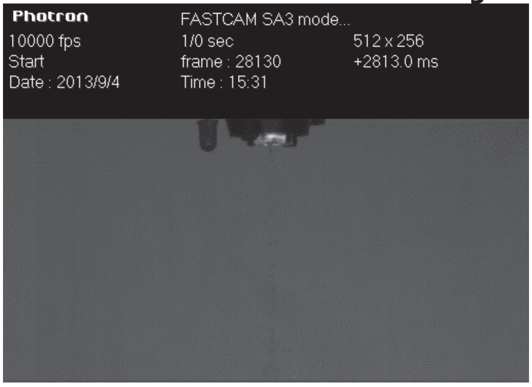

Figure 3.2: Completed deactivation of a solenoid valve exhibiting residual pressure build-up. Notice the dribble.

Further investigation revealed that nozzle bodies with a larger reservoir behind the nozzle generated residual pressure in the nozzle body. The residual pressure delayed the cut off of the nozzle and caused a noticeable slow in flow before stopping the flow, as shown in Figures 3.3 and 3.4. This implies that a suitable nozzle body needs to have a small reservoir for the fluid flow behind the nozzle to enable direct fluid cut-off and also to alleviate dripping when not spraying.

Table **3**3 displays the ‘turn on’ and ‘turn off’ results for the four commonly available solenoid valve spray nozzles. It can be seen that the Asco valve has the fastest turn on times (between 5.8 ms to 6.1 ms) and the second fastest turn off times (12.2 ms to 12.4 ms). The Weedit solenoid valve had the fastest turn off time of between 10.9 ms to 11.2 ms, approximately 1.2 ms faster than the Asco valve.

Table 3.3 Time for solenoid activation and deactivation at 40 psi. Timing for activation is measured from when the LED illuminates to when fluid first appears from the nozzle. Deactivation is measured from the when the LED turns off to when the fluid disperses from the nozzle.

	Switch On				Switch Off				
Solenoid	Run One Time (ms)	Run Two Time (ms)	Run Three Time (ms)	Average Time (ms)	Run One Time (ms)	Run Two Time (ms)	Run Three Time (ms)	Average Time (ms)	State of Spray
Goyen	11.2	10.3	10.3	10.6	21.9	20.7	21.0	21.2	Cut off
Capstan	8.2	8.1	8.1	8.1	15.3	12.8	13.3	13.8	Collapsed
Weedit	8.8	8.6	8.6	8.7	10.9	11.2	10.9	11.0	Cut off
Asco	6.1	6	5.8	6.0	12.6	12.2	12.4	12.4	Cut off

The electrical operation of the solenoid was important to assess when investigating the optimisation of the solenoid ‘turn on’ and ‘turn off ‘ times. When a solenoid is turned on, the coil resists the increase in current flow and creates a ramped up response to activation. This ramped up response to activation is the delay that is seen at the ‘turn on’ of the solenoid. One way to enhance (reduce) the ‘turn on time’ is to develop a trigger circuit for the solenoid which increases the voltage above its holding voltage (termed ‘over driving’) and then reduced to its holding voltage once fully activated. Upon turning off the solenoid, the coil again resists change and must dissipate the current in it. It does this by applying a reverse voltage spike on the supply line. The spike must be dissipated as quickly as possible to allow the coils magnetic field to collapse quickly.

The NCEA has designed a circuit that implements the above requirements to minimise the ‘turn on’ and ‘turn off’ times.

Recommended spray delivery system and design

The research has shown that the optimal system for the spray delivery will comprise of a narrow angle nozzle of 15° for uniformity of the spray fan delay time. A manufactured nozzle body with a small reservoir for the fluid to pass through between the solenoid valve and the nozzle is required. This will reduce the residual pressure in the nozzle body which causes dripping and a non-uniform fluid cut-off.

A further feature to optimise the time delay is to implement a trigger circuit which will ‘over drive’ the turn on and quickly dissipate the turn off voltages. Asco valves have been selected for these optimisations because Asco valves are readily available and have the overall best on time and off time of the commonly available spray nozzles. The optimisations will further enhance the performance of the Asco valves.

Between the rows weed control

Three equipment options have been developed for trials of between the row weed control, which are herbicide application, tillage and a combination of both. Herbicide and non-herbicide (i.e. tillage) options will be trialed to evaluate which is most effective for weed control with the least damage to the pyrethrum plant.

Between the row weed control is different to in the row weed control. This is because there is not crop between the row and the operator is capable of using different weed control options to kill the weeds in this crop free zone.

Herbicide treatment can be applied unshielded to the crop free zone by a spray applicator positioned above the crop free zone. This method has shown significant issues for drift in other row crops which would not be an issue when using a selective herbicide that the pyrethrum is resistant to. However, the selective herbicides are no longer obtaining the same level of control as they once were in the pyrethrum and different herbicides need to be used to obtain a satisfactory weed kill.

A between the row shield can be used to allow herbicides that the pyrethrum is not resistant to be used. The shield, ‘shields’ the pyrethrum from the herbicide and ideally contains the spray and drift inside the shield. There are three significant issues to address in using a shield spray system in pyrethrum which are:

Drift: The shields create an artificial area inside them and as the shields pass over the ground they can sometimes create turbulence within the shield, which can cause the drift to escape from under the shield.

Weed coverage: This is where the spray pattern inside the shield should cover the ground area under the shield but not the side of the shields. This is because if the herbicide is sprayed onto the side of the shield, it can act like a wick wiper onto the crop.

Accidental crop application: This is where the crop sometimes has larger pyrethrum plants which grow into the crop free zone and as the shield passes over these plants it sprays them with herbicide, which can be translocated back to the root zone and damage the pyrethrum plant.

To assess shield spraying in the pyrethrum industry and the three issues outlined above, a shield has been purchased from Micron Technologies and mounted onto a depth wheel parallelogram to maintain a constant height. Additionally a set of cutting discs has been added as an option to the parallelogram. The cutting discs will cut off the parts of the pyrethrum that grow into the crop free zone. Figure 11 shows a ‘stitched’ depth image of pyrethrum on the left and the white lines at either side indicate where the shields will be positioned. This image shows that some of the pyrethrum plants exceed this width and cross the white lines. This is the portion of the plant that needs to be ‘trimmed’ to stop translocation of the herbicide to the roots. Further a variety of different nozzles can be fitted to the shield to determine the best performance for drift and coverage.

Figure 3.5 is an engineering drawing of the parallelogram with cutting discs to cut the pyrethrum growing into the row and a shield spray unit at the rear.

A second means of weed control for weeds in the crop free zone is inter-row cultivation. This requires a piece of tillage to be used between the rows and cut the roots of the weeds via a ‘point’ or ‘sweep’. To evaluate this, a parallelogram has been fitted with a single cutting disc in the centre and a sweep at the rear. The cutting disk is required to cut and large lumps of stubble or vines and stop them wrapping around the tyne. Figure 3.6 shows the parallelogram with the cutting disc at the front. The cutting disc also acts as a depth wheel keeping the sweep at a constant height as well as making sure nothing wraps around the tyne. The tyne is the long, black vertical bar at the rear of the parallelogram and the sweep is the green ground engaging tool at the bottom of the tyne

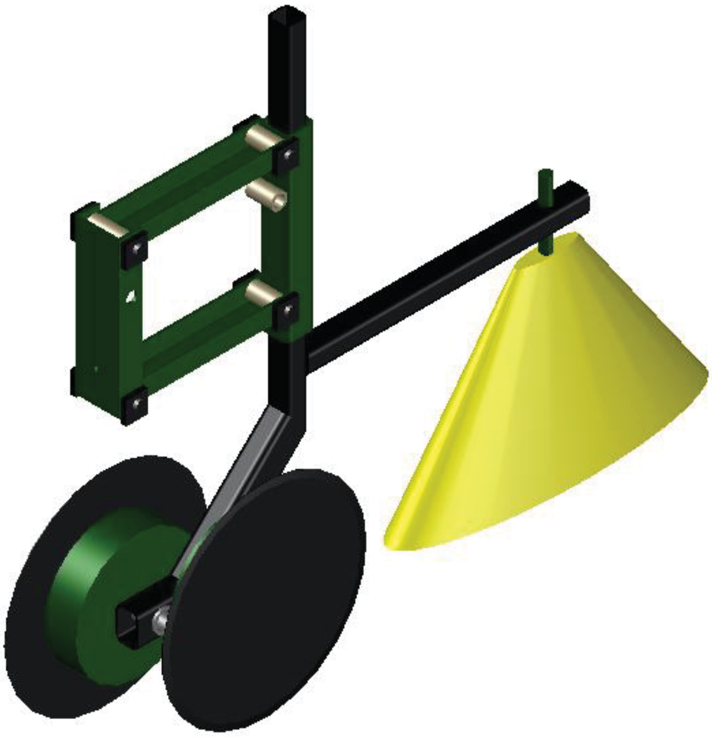

Figure 3.5: Parallelogram with cutting discs and shield.

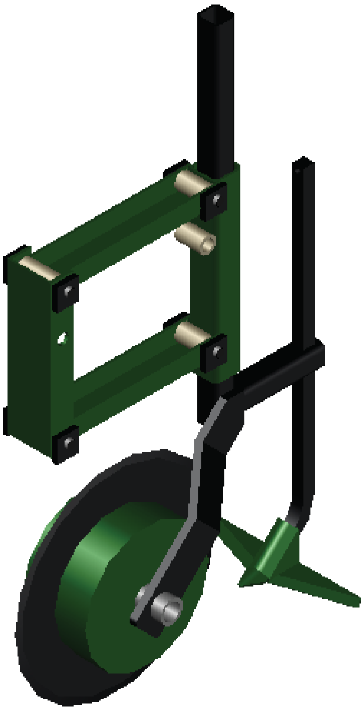

Figure 3.6: Parallelogram with single cutting disc and tyne and sweep assembly.

Guidance

Guidance is necessary for in crop weed control to ensure that the weed control treatments are consistently applied on target (i.e. between the rows for shield spraying and on the row for spot spraying). Two guidance systems will be evaluated for the Milestone 105 report. One will be a commercial system supplied by BRA and a second is a guidance system that is developed using the image information gathered by the weed spot spray software developed in this project.

In order to provide adequate between the row weed control, a guidance system will need to be used to keep the spray shield and tillage implement centred between the rows of pyrethrum. The machine vision system developed in this project is used for in row weed control has been further developed to allow for implement control as well.

The further development has been to 'stitch' the depth images and the RGB images together so that a history of where the system has been is recorded (Figure 3.8). The pyrethrum is identified in both the depth image and the RGB image and a centre line is fitted to the RGB image and this centre line will then control a side shift hitch to align the implement in the centre of the crop line. Figure 3.8 shows the stitched depth and RGB images. The depth image displays the height of the plants in different colours. Blue being the lowest and the warmer the colour, the higher it is. The white lines indicate the position that the shields will be in. The RGB image shows the 'stitched' colour image of the pyrethrum and the white lines indicates the same as for the depth image. The pink line is where the system believes the centre of the crop is and can then send a signal to the side shift hitch to relocate. From visual inspection it can be seen that the pink line follows the centre of the crop reasonably well.

A side shift hitch is a three point linkage quick hitch that attaches to the tractor at the front and to the implement at the rear. These particular quick hitches are unique in that they also have a hydraulic ram built into them. The rear is coupled to the front via large roller bearings and the ram can push and pull the rear section sideways relative to the front section. Figure3.7 is a visualisation of this implement.

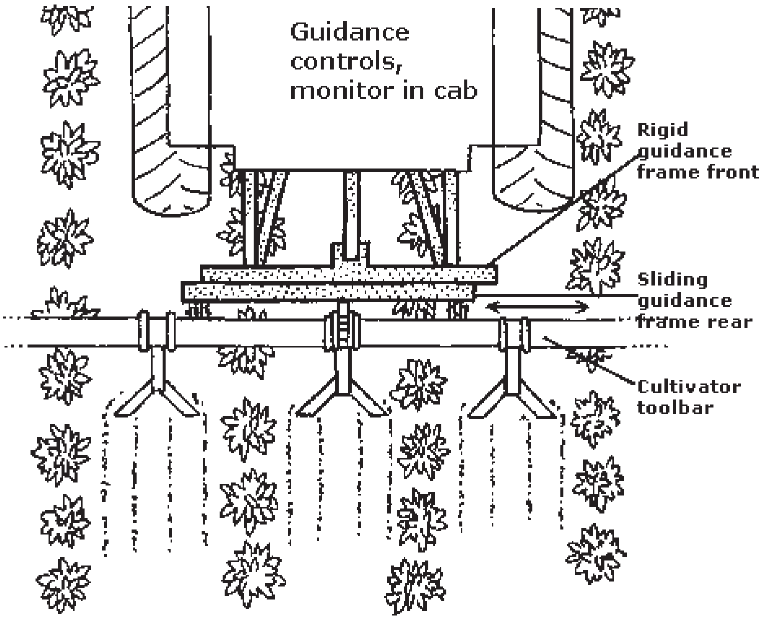

Figure 3.7: Drawing of how a side shift hitch operates reproduced from http://www.sare.org

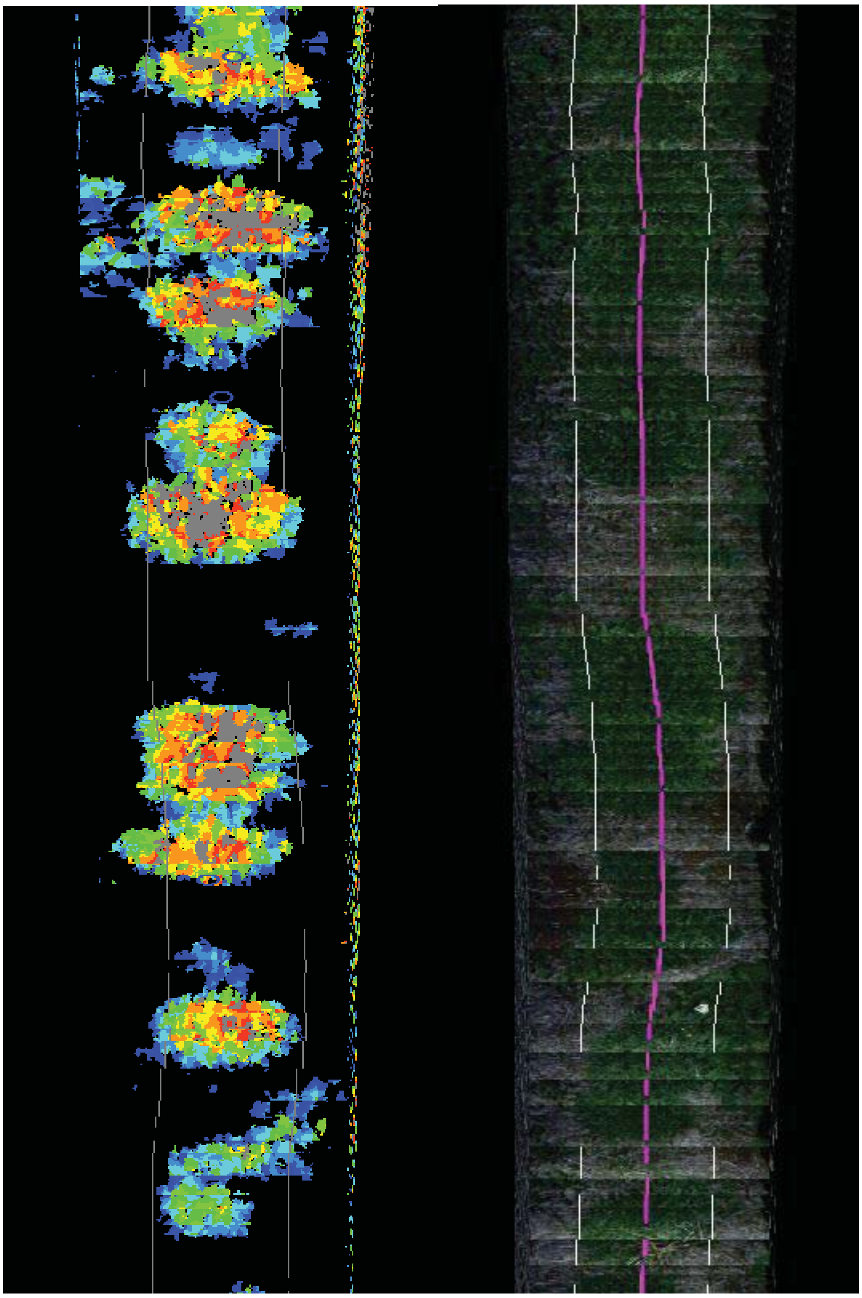

Figure 3.8: Stitched images of depth and RGB showing the determined centre line in pink and the in row area between the white lines.

Refine image analysis algorithms

The refinement of the image analysis algorithms is an ongoing process. To date additional work with the algorithm has seen the combination of statistical depth features and Local Binary Pattern (LBP) features, used with a neural network classification function. The results from this increased the accuracy of the LBP from 57.9% hit rate (from Milestone 103) to 90%. This is a very encouraging result and at a field ready level however ongoing work will see the depth-LBP features combined with the custom algorithm (from Milestone 103) to increase the performance of the combined classification.

Assemble machine vision equipment for the field ready prototype

The field ready prototype has been manufactured by Scott Cox engineering in three distinct sections. The sections are the frame, the light diffusing hood and the attachments for the spray shield and tillage options. Figure 3.9 shows a 3D cad drawing of the components as well as photos of the manufactured and assembled frame and Hood. Figure 3.10 also displays a photo of the camera and lighting module that was used on the data collection unit in Milestone 102 and Milestone 103; it will also be used in Milestone 104 and 105.

Field prototype

In consultation with BRA it was decided that a field ready prototype should be developed which could be used in Milestone 105 as well as be retooled for use in Milestone 106. The developed field prototype can be operated in the 60 cm row spacing of 'speedling' plantings as well as the 40 cm row spacing in the 2013 plantings and covers a full bed width of 1.8 m. The field ready prototype has been designed so that it can operate in the 'speedling' rows of pyrethrum which are 70 cm apart as well as over the more narrow rows which are 40 cm apart on a 1.8 m bed.

The prototype has a platform off to each side to hold the spray apparatus and also wash water. It maintains a camera height of one metre and can mount a parallelogram on the rear to operate between the rows. Figure 3.9 is a 3D model of the assembled field ready prototype produced from engineering drawings developed at NCEA. The prototype is intended to house between one and three machine vision units on the top of the black internal frame of Figure 1 dependent upon the row spacing of the pyrethrum. Figure 3.10 contains a photograph of a single unit of the machine vision camera system. The computer for the system is positioned in the motor vehicle and is the same unit as was used in Milestone 103 report for the data collection. Figure 3.11 is the manufactured and assembled field ready prototype.

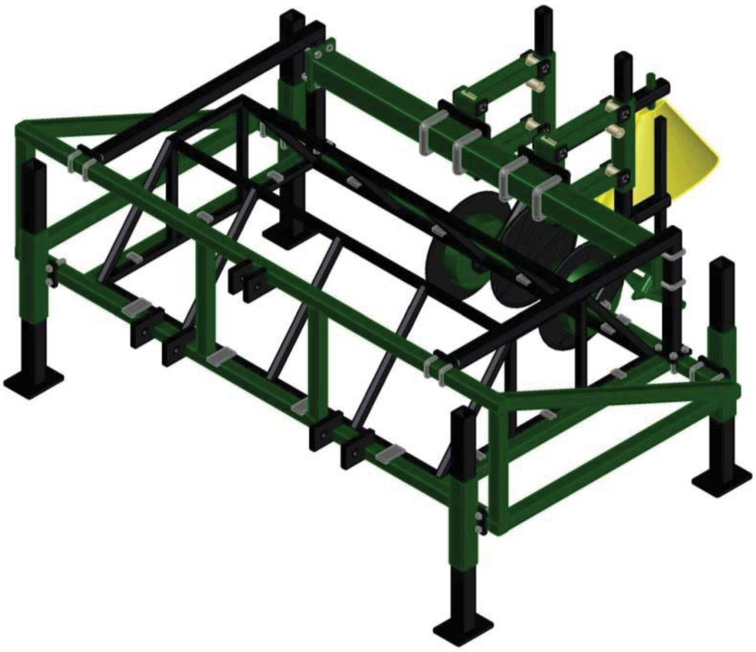

Figure 3.9: 3D model of prototype field unit based on NCEA's engineering drawings.

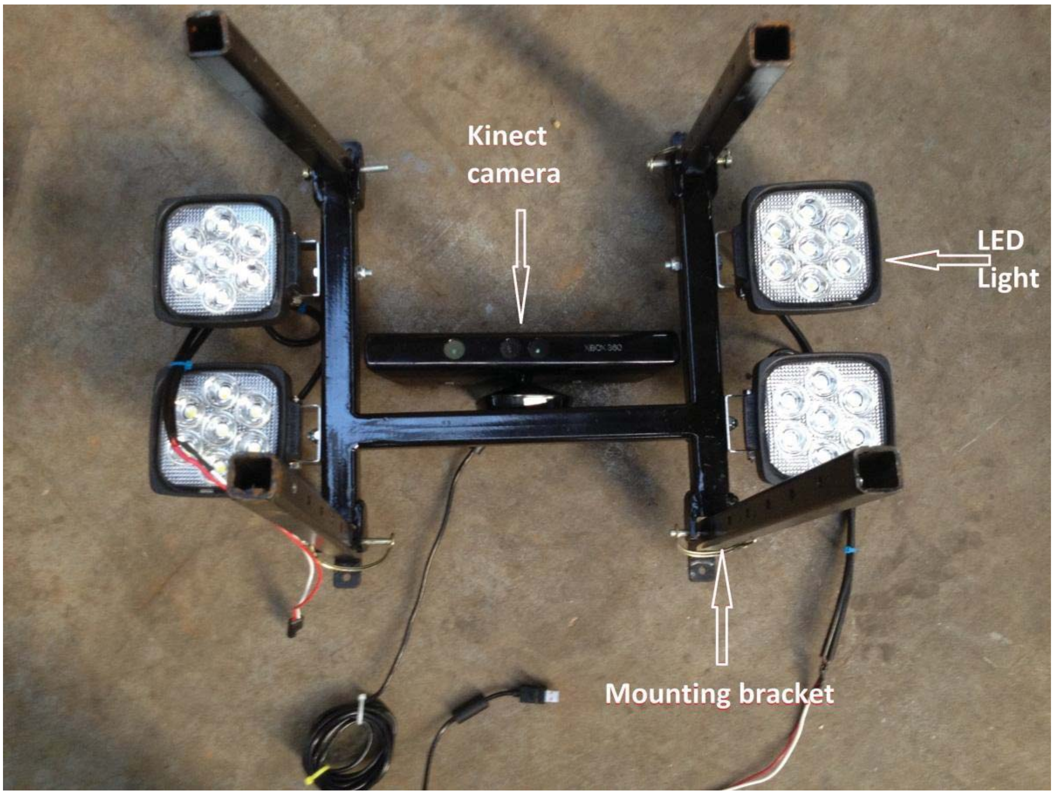

Figure 3.10: Camera and lighting module used previously in Milestone 102 and 103 and will be used in Milestone 104 and 105. Each unit can analyse one metre width on the ground.

Figure 3.11: The manufactured and assembled field ready prototype frame.

Milestone 105

The project has been following the milestones satisfactorily. This milestone has been focused on deploying the prototype spot spray and inter row machine effectively in the field in actual pyrethrum crops. The prototype system incorporates a vision guidance system, a vision spot spray system for weeds in the row and the capability of inter row shield spraying and cultivation. The vision guidance system is for maintaining the spray nozzle position precisely where they are required to be to supply adequate dosing to the weeds in the row and avoid overspray to the pyrethrum plants as well as positioning the spot spray machine vision technology and the inter row weed spray shields and cultivating gangs. The guidance system has been tested with an effective accuracy of plus/minus 5cm which is satisfactory for the purposes required.

Shield spraying has been conducted on several BRA sites with effective results and appears a very promising method for inter row weed control. Inter row cultivation also tested well and longer term evaluation is required to determine if the soil disturbance is an issue for weed seed germination. The technologies developed for the timing and minimisation of the distribution footprint of the herbicide tested satisfactorily at plus 2.5cm each end of the weed. The testing of the individual developed components has proved satisfactory to allow integration into the phase which is a 1.8 m prototype machine

Evaluate the precision spray delivery system, spray shields and guidance system

Prototype setup

The prototype machine outlined in milestone report 104 was transported to the BRA offices at Ulverstone and deployed to operate on a Garford side-shift hitch. To fit the prototype to the side-shift hitch, the prototype had the top link position of the three point linkage modified to suit the side-shift hitch as the side shift hitch position was non-standard. All other three-point-linkages positions matched.

The prototype machine outlined in milestone report 104 was packed up and mounted onto a pallet for transport. Road transport was used to move the prototype machine to the BRA offices at Ulverstone, Tasmania. The prototype machine had the Garford side shift hitch mounted onto the three-point-linkage pins to provide a means of correcting the prototypes position in the row. Modifications were made to the top three-point-linkage link as the Garford hitch had non-standard spacing. The modifications were carried out at the BRA staff at the BRA maintenance shop in Ulverstone. The components of the prototype were all re-assembled onto it and the complete prototype and hitch was mounted onto a BRA tractor for testing.

Figure 4.1: Prototype machine and Garford side-shift hitch mounted on the tractor, shown in a crop of pyrethtrum.

Guidance testing

A vision guidance technique had been developed and was tested on 40 cm row widths of pyrethrum. The technique proved effective in all situations with an accuracy of plus/minus 5cm. To test for accuracy the prototype system was run as in Figure 4.1 with the coulter disk for the tillage apparatus (at the rear of the machine) leaving a line on the ground. The guidance system, via the side shift hitch (Figures 4.2 and 4.3) moved the prototype left and right to the direction of travel in an effort to keep it centred on the row. Therefore accuracy can be determined by how far off centre the coulter disk line is. A measuring stick (Figure 4.4) was manufactured that allowed the user to walk along the row of pyrethrum holding the stick over the pyrethrum being able to see the variation in the lines position. These results were based on several kilometres of operation on varying sized pyrethrum.

Tillage testing

A tyne with a 20 cm sweep was trialed with satisfactory results. The guidance system was able to locate the tyne correctly so that no pyrethrum was displaced and the sweep was able to go under the overhang of the larger pyrethrum without causing any damage. Where the pyrethrum was small there was the occasional plant that had soil thrown on it. The areas that were cultivated were marked and evaluated after a week and after the next rain event to evaluate the weed kill from the tillage and also the weed seed germination caused by the ground disturbance. The weed kill was effective and weed seed germination appeared no different to the shield sprayed area.

A tillage tyne with a 20 cm sweep operating between the rows of pyrethrum for weed control was trialled with satisfactory results (Figure 4.5). The guidance system was able to locate the tyne assembly centrally in between the rows (plus/minus 5cm) so that no pyrethrum was cut displaced or damaged when operating and the sweep was able to go under the overhangs of the larger pyrethrum plants without getting caught on them and causing any damage to the pyrethrum plant. Where the pyrethrum was small there was an occasional plant that had soil thrown on it however, the test team did not think this was a significant issue. The test rows that were cultivated were marked with a flag and evaluated after a week and after the next rain event to evaluate the weed kill from the tillage and also the weed seed germination caused by the ground disturbance by the tyne assembly. The weed kill was 100% effective for the weeds that fall within the 20cm band of the sweep and weed seed germination in the tilled rows appeared no different to the rows that were shield sprayed only.

Figure 4.2: Side shift hitch in the fully left position

Figure 4.3: Side shift hitch in the central position

Figure 4.4 Measuring stick used for measuring accuracy of the guidance system

Figure 4.5: Image of the tyne assembly in operation. The coulter disk in the front of the tyne cut any vines and trash so that it could not wrap around the tyne.

Inter-row shield spray test

The inter-row shield spray prototype was manufactured to operate on the 60 cm rows and was unable to be operated as intended due to not fitting down the rows. BRA maintenance-fabrication has modified the inter-row shield spray unit to fit into 40 cm rows.

Shield spraying has been undertaken across several BRA sites with acceptable results. It is difficult to evaluate inter row shield spraying with hit and miss results as it is not feasible to individually evaluate each weed between the rows and it is difficult to ‘see’ the individual effects of drift on a pyrethrum plant for a statistical approach but it is possible to notice overall effects.

‘Spray footprint’ speed timing test

A test was set up on the roadway into the ‘Werrin’ site. Speeds between two and seven km/h were tested with the start of spray fluid matching the start of the plant on all test runs and the spray fluid stopping within 2.5 cm of the end of the plant in all test runs.

A test was set up on the roadway into the ‘Werrin’ site. Speeds between 2 and 7 km/h were tested with the start of spray fluid matching the start of the plant on all test runs and the spray fluid stopping within 2.5 cm of the end of the plant in all test runs. The test consisted of repositioning six compact weeds onto the gravel road surface 20 meters apart from each other in a line. The prototype was then calibrated to ‘trigger’ off the weeds and blue dye was added to the spray water. The system was then run over the line of weeds at speeds of 3, 5 and 7 km/h. Each speed was held constant for the full run. The spray on and off positions were evaluated for each speed by visually viewing the blue dye spray pattern on the ground and over the weeds. New weeds would be repositioned to unsprayed areas of the road for the next test run. Figure 5.5 shows the spray nozzles spraying blue dye onto a weed in a regrowth pyrethrum crop.

Figure 4.5: Spray nozzles spraying blue dye onto a weed in a regrowth pyrethrum crop.

Evaluate weed detection under a range of conditions

Evaluation of the weed detection algorithm at the ‘Werrin’ site near Ulverstone occurred over April and May 2014. The evaluation of the algorithm showed high accuracy on the small weeds and large weeds in the pyrethrum crop with high accuracy in identifying pyrethrum like plants that were over the entire pyrethrum size spectrum.

Evaluation of the weed detection algorithm was carried out at the ‘Werrin’ site near Ulverstone, Tasmania. Occurring over April and May 2014. The evaluation of the algorithm showed high accuracy on the small weeds and large weeds in the pyrethrum crop and high accuracy in identifying pyrethrum ‘like’ plants over the entire pyrethrum size spectrum i.e. from 30 cm width down to 10 cm width. The testing was achieved by using blue dye in the spray water and visually determining where the system had sprayed. It was difficult to get percentage ‘hit’ results as the weed growth in the crop was at the out of hand stage and manual chipping was being carried out to regain control of the weeds and crop. The team was able to determine an overall satisfactory result for the small and large weeds and isolate areas for further refinement of the algorithm to be tested next season.

Figure 4.6: Blue dye on weeds in a pyrethrum crop.

Modify machine vision hardware as necessary

No modification was necessary.

Milestone 106

Mid Term Project Review of: PY12000 Precision weed sensing for pyrethrum

Held via teleconference on: 25/9/2014

Background

It is Horticulture Australia Limited (HAL) policy that a mid-term project review is required for all projects with a life of project value greater than five hundred thousand dollars ($500,000), or a project life of three years or more. A mid-term project review ensures that a project is able to document and report on how it is progressing in terms of delivery of outputs and achieving intended outcomes. Recommendations from the mid-term project review shall be incorporated as appropriate into the remainder of the project, considered in future milestone reporting requirements and may also be utilised in the development of future projects. Benefits to industry will be greatest when projects are reviewed and project methodology is refined to meet the needs of industry. This project is subject to the outcomes of this review, and may include modification and/or termination.

Process

The mid-term was in the form of an interview with key project personnel undertaken by HAL via teleconference facility.

Key project personnel included:

- Cheryl McCarthy
- Tim Groom
- Craig Baillie
- Steven Rees

Input to the review was provided by Anthony Kachenko, HAL Team Leader R&D.

Objectives

HAL determined the objectives for the mid-term project review were as follows:

1. Determine that this project was being managed responsibly and that all contractual requirements with HAL are being adequately met.

2. Review activities undertaken during the study including suitability of the methodology and benefits to industry.

3. Assess the quality of outputs and the level of engagement with stakeholders and identify how/if engagement could be enhanced.

4. Review commercialisation opportunities and identify potential impediments to effective development and commercialisation

5. Undertake a SWOT analysis.

6. Determine recommendations and agree on activities for the remainder of the project.

Interview Questions

The following questions were used as the basis for the mid-term review.

1. How would you rate the value of this project to the pyrethrum industry?

   Not at all important 0 1 2 3 4 5 6 7 8 9 10 Very Important (please circle)

   The review team rated the project 9/10.

2. Overall, how would you rate your level of satisfaction with project activities in terms of addressing expected program outcomes? Key project outcome is *performance specification and usage guidelines based on the proof-of-concept technology for automated spot spray weed control strategies.*

   Not at all satisfied 0 1 2 3 4 5 6 7 8 9 10 Very Satisfied (please circle)

   The review team rated the project 7/8. Milestones have been met as per the project plan. Testing of the prototype spot spray has proven successful, pending modifications to accommodate row spacing.

3. Please comment on how well you think the program is delivering outcomes and benefits to stakeholders. [*May prompt in the key areas of outcomes*]

   Program is on track but still to be proven in the field. Novel IP has been generated with NCEA recently lodging a provisional patent in the United States for precision weed spray technologies. An ongoing relationship between NCEA and John Deere regarding the spot spray technology was noted.

4. In your opinion, what are the strengths, weaknesses, opportunities and threats to this project?

   The following SWOT analysis table were compiled by the review team.

Strengths	Weaknesses
1. USQ is very competent, dedicated and determined to make the project work. Strong working relationship noted. 2. USQ experienced in this area with a sound track record. 3. BRA is in control of the technology and is able to implement trials on its own farms – it is not reliant on other people to do the field testing. 4. BRA is ready to build a commercial unit as soon as the project demonstrates that the concept works.	1. Location of USQ in Toowoomba removed from the pyrethrum growing area in Tasmania. 2. The field trials have relied on the BRA workshop to make some modifications to the machinery, which is not their main pursuit and has created some delays. 3. Duplication with other technologies. 4. Level of adoption and window of adoption.

Opportunities	Threats
1. There are opportunities in a range of horticultural crops, for example the potential to overcome volunteer self sown potatoes. 2. Our initial appraisal suggests that this technology could halve the current cost of weed control, improve post-harvest crop health and control difficult weeds 3. Non-herbicidal control methods warrant reviews as part of the machine's capability e.g. steam/foam, mechanical means etc.	1. The project is at the proof of concept stage, and whilst good progress has been made on weed versus pyrethrum recognition, the technology has not been confirmed that the technology has a practical place in commercial production.

5. From your perspective, what is not working well or what areas would you suggest changes are needed in the future?

   The committee were of the opinion that the project was working well and no changes were warranted.

6. Any closing comments or suggestions for the future of this project?

   Anthony advised that a HAL funding call was to open on 29 September and would include opportunities to submit proposals surrounding mechanisation and in field robotics. He advised the committee to review the HAL website for further information.

   Craig noted that a forum in Canberra was being held on October 10 to discuss robotics with key RDCs and players working in this environment. He noted that HAL should attend.

Recommendations

PY12000 to continue as per the contracted milestones.

Milestone 107

Progress since the Milestone 105 report (July 2014) and Stop/Go Milestone 106 (September 2014) has focussed primarily on algorithm refinement to specifically target weeds that are similar in overall geometry to pyrethrum. The algorithm refinement is showing promising results and will be validated in field trials in 2015.

The field trials will use the 1.8 m multi-row field ready prototype that was originally developed in Milestone 104 and that was determined in Milestone 105 to be suitable for subsequent field evaluations. In consultation with BRA, the field ready prototype from Milestone 104 was constructed to enable its operation for field scale evaluations, in lieu of rebuilding the field ready prototype in Milestone 107.

Construct 1.8m multi-row field ready prototype for commercial testing

Milestone 104 developed a 1.8m field ready prototype (report in November 2013) which was determined in Milestone 105 (report in July 2014) to be suitable for subsequent field evaluations for the project. BRA funded USQ an additional $10,000 to top up the proof of concept production costs, external to the project funds to produce the field ready pre-production prototype. Figure 1 shows the field prototype attached to a BRA tractor at Werrin.

Figure 5.1: The pre-production prototype attached to a BRA tractor in a three month old crop of pyrethrum on a BRA leased farm (Werrin).

Collate weed detection results

Algorithm development and testing has been ongoing. Algorithms that were reported in Milestone 105 performed satisfactorily for weeds that were smaller or larger than pyrethrum. Refinements to the plant discrimination algorithm have been trialled to identify weeds that were similar to pyrethrum in shape and size. Figure 5.2 shows the field outcome of identifying weeds that were smaller than the pyrethrum in the evaluation outlined in Milestone 105 report. The small weeds have been sprayed with blue dye in the image. They weed may still have been occluded (i.e. touching other plants) but the image analysis isolated the individual plant and size and minimum height was that able to identify it as a weed.

Figure 5.2 showing the weeds correctly identified and targeted with blue dye.

Data collection

The data was collected by Francis Chamley of BRA from post-harvest regrowth at the trial site on Greg Gibson's property at Hagley, Tasmania. The data was collected at two different times. The growth stage of the pyrethrum was as it was coming out of its winter dormancy and starting to grow. The first recording of 30.8 GB was on 28 August 2014 and the second collection of 90.4 GB of data on 25 September 2014. The data collection was obtained from the pre-production prototype unit and therefore is real world data. i.e. the same as will be encountered in the field in 2015.

Algorithm

The operation of the complete image processing system for the identification of pyrethrum from weed for spot spraying will apply the techniques previously reported (Milestone 105 report) to highlight the particular plant to be identified and to determine if it is smaller or larger than the pyrethrum in the field at that particular growth stage/s.

Figure 5.3 shows an image of the analysis for identifying individual plants by finding the pinch points. A pinch point is where the width of plant is at its narrowest. If a pinch point is found on both sides of the plant then it is deemed the point at which one plant ends and the next plant begins. We can see if Figure 5.3 the pinch point is located on both side of the plant in the left hand image (denoted by red circles) and that this point is pointed to in the colour image of the right, highlighting satisfactory isolation of plants from each other for analysis. Other points are shown with the red circle but they are not on both sides of the plant which is required to determine plant end.

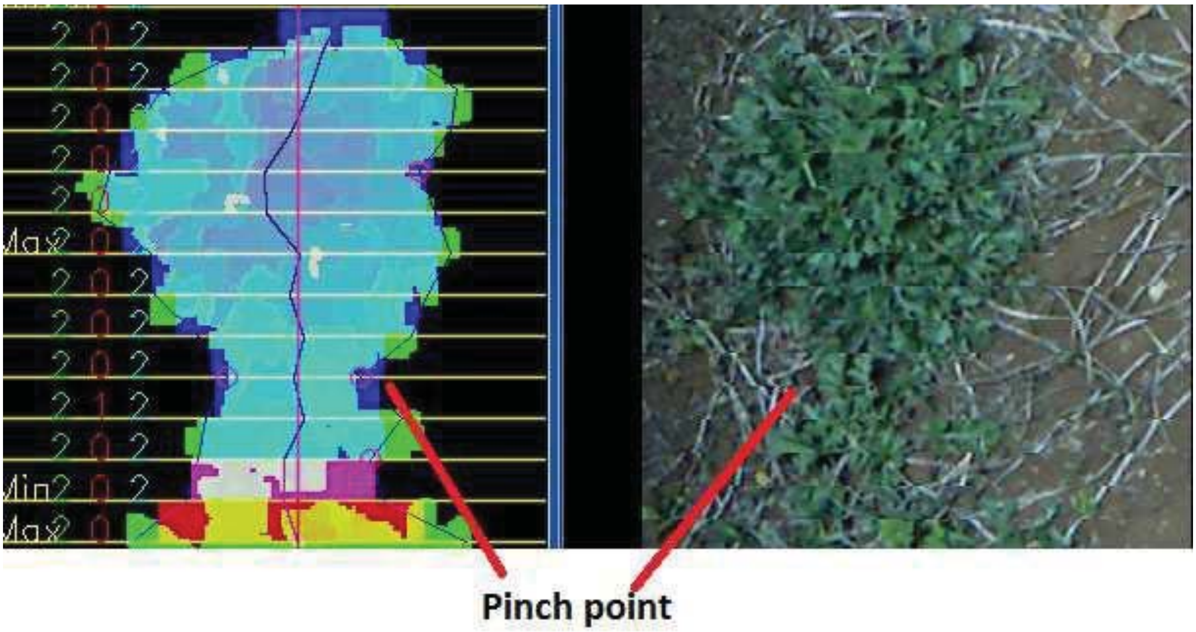

Figure 5.3 The image on the left is the analysis output of determining the pinch point of the plant by a red circle on both sides. The pinch point is highlighted on the colour image on the right hand side.

Milestone 108

Progress since the Milestone 107 report (November 2014) has focussed primarily on in-field evaluation of the spot spray system. The field evaluation used the 1.8 m multi-row prototype that was originally developed in Milestone 104 and determined to be suitable for field evaluations in Milestone 105. Further refinements will be performed in preparation for planned commercial-scale field trials in September.

Refine the prototype

Modifications to the prototype were undertaken to enable the 1.8 m multi-row prototype to spot-spray each individual row in a bed of four rows. Modifications included wiring, mounting and testing the two quad-core computers responsible for controlling a camera and four solenoids. Division of the camera image into left and right regions of interest facilitated two analysis streams for automatic detection and control of the solenoids. A power control box for the lights, camera, computers, pump and solenoids was installed on the rear of the 1.8 m multi-row prototype. The 1.8 m multi-row prototype was connected to a control computer running a user interface, allowing the operator to configure and instruct the 1.8 m multi-row prototype via a CAN communications link from within the tractor cab. Power was supplied to the 1.8 m multi-row prototype from the battery of the tractor.

In February 2015, Bruen Smith and Matthew Tscharke travelled to Ulverstone to fit out the 1.8 m multi-row prototype. The 1.8 m multi-row prototype consists of two detect-spray systems, a control computer and a control box to allow simultaneous weed detection and spraying of a bed of four rows. The detect-spray systems each monitor two adjacent crop rows and each consists of:

- one camera which views two adjacent crop rows;
- two pairs of spray solenoids (hence, four spray solenoids for each detect-spray system), with each solenoid covering half of one crop row; and
- one multi-core computer to analyse the images from the camera, trigger the spray solenoids and communicate with the control computer in the tractor cab.

A control computer with user interface is positioned in the tractor cab, and communicates with each detect-spray system via a CAN communications link. The control computer allows the user to configure software parameters of the detect-spray systems. A control box with master electronic switches for the lights, camera, computers and pump is positioned on the prototype (Figure 6.1). Figure 6.2 shows the 1.8 m multi-row prototype, which is normally towed by the tractor, but here shown raised off the ground for portability after a data collection event.

Figure 6.1: The control box (*left*) and case containing the two detect-spray computers (*right*).

Figure 6.2: The 1.8 m multi-row prototype after a data collection event.

Evaluate and refine weed detection

Data totalling 95 GB was captured on two occasions from a trial site located at Kindred, Tasmania in early 2015. Data comprised pyrethrum crop after harvest and after mulching, pyrethrum flowering, weeds amongst pyrethrum and pyrethrum regrowth. Collected data was passed through the image analysis algorithm to simulate on-farm operation and determine appropriate settings for real-time field evaluation of the 1.8 m multi-row prototype.

The first recording of 55 GB of data of pyrethrum was collected in February 2015 after harvest and mulching, when the pyrethrum was in regrowth and weeds were not abundant. A further 40 GB of data was collected in April 2015 after regrowth had occurred and weeds were more abundant. The data sets contained pyrethrum after mulching (February 2015), pyrethrum regrowth after mulching (April 2015) and pyrethrum flowering and weeds amongst pyrethrum (also April 2015). Variation in plant appearance for these stages is shown in Figure 3.

The data collected with the 1.8 m multi-row prototype was stored in an electronic format that enabled subsequent desktop algorithm development and evaluation, and also enabled post-processing of the stored data to emulate streaming and processing of real-time data in the field and generation of spray solenoid trigger signals. Hence, algorithm parameters that were optimised in desktop evaluations could be ported directly to the online prototype for real-time field evaluations.

a) b) c)

Figure 3: Pyrethrum at different growth stages: a) mulched pyrethrum with a small amount of regrowth; b) regrowth after mulching c) regrowth and flowering

Evaluate the pre-production field prototype under commercial conditions

The field site selected for real-time evaluations had the required crop and weed conditions but non-typical row configuration which necessitated modification of the 1.8 m multi-row prototype. The 1.8 m multi-row prototype was modified to detect weeds only within the two centre rows of a four-row bed, which was a physical modification that did not alter operation of the weed detection algorithms.

Evaluation of the 1.8 m multi-row prototype's weed detection ability within the crop row was achieved by towing the prototype by tractor at approximately 3 km/h over a total of 242 pyrethrum plants. The 1.8 m multi-row prototype was set up to generate spray patterns that could be easily visually discerned and achieve good row coverage when a weed was detected. Spot spray hit rates for the weeds (true positives) and pyrethrum plants (false positives) were recorded. The overall weed hit rate was evaluated to be 75% and is considered an exceptional result because:

- Weeds varied greatly in appearance, ranging in size from less than 1 cm to greater than 30 cm in diameter, and often overlapped with the pyrethrum plant edges and one another. Further analysis will determine the distribution of detection rate versus weed size.
- 5-10% of misses were due to the spray solenoids rapidly switching and not turning on in time (i.e. the weed was detected but not precisely sprayed – this will be addressed by finetuning the spray solenoid parameters).
- 5% of misses were too small to visually discriminate as pyrethrum or weed.

From the evaluation, 88% of the pyrethrum plants were correctly identified as pyrethrum and not sprayed. The remaining 12% were false positives. Half of the false positives were applied with non-lethal doses, since the false positives occurred over edges and sections of pyrethrum clumps. The other half of the false positives occurred over pyrethrum that had delayed regrowth with noticeably different appearance. Incorporating data of pyrethrum plants with abnormal appearance into the training set is expected to reduce the number of false positives and will be verified in the coming months by additional datasets collected during extensive trials planned for September.

Camera and nozzle configuration for spray application evaluation

Row beds were identified on a trial site at Kindred, Tasmania to perform real-time field evaluation of the 1.8 m multi-row prototype in May 2015. Pyrethrum plants on the two outer rows at the trial site were observed to be falling off the bed into the wheel track, due to the specific practices adopted at the trial site, and were not considered to be typical conditions. Therefore, the 1.8 m multi-row prototype was modified to analyse the centre two rows only during the real-time field evaluation, which enabled the algorithm performance to be evaluated on crop rows that were typical of standard commercial conditions.

Two regions of interest (ROIs), over the two centre rows of the bed, were manually configured for real-time field evaluation of the 1.8 m multi-row prototype. The operation of the algorithm was evaluated without the guidance hitch (reported in Milestone 107), as the driver was observed to follow the row with sufficient accuracy for the purpose of algorithm evaluation. ‘Sampling error’ occurred where the operator did not align the ROIs over the crop row and the ROI drifted into the adjacent row (Figure 6.4). Implementation of the vision guidance (reported in Milestone 107) will assist with maintaining the ROIs over the row.

a) **b)**

Figure 6.4 : a) left and right ROI; b) example of sampling error (i.e. ROI not over the row - guidance will overcome this problem)

Nozzle selection for spray application evaluation

A pair of solenoids corresponding to each ROI was installed on the 1.8 m multi-row prototype (hence, four in total). Use of a pair of solenoids enabled the position and coverage of the spray application to be determined during field evaluations. The selected nozzle tips were as follows and are displayed in Figure 6.5:

- A nozzle next to the row, to indicate spray position: 15° nozzle with flow rate of 4 gallons per minute capacity (GPM) at 40 PSI (TeeJet 1504); and
- A nozzle over the row, to indicate spray coverage: 0.25 GPM at 40 PSI (UniJet 150025).

The TeeJet 1504 tips were higher flow rate and delivered substantially more dye-water mixture, hence the applied spray formed a highly-visible spot on the ground that was easy to identify at eye level. When a larger weed was detected, a series of spots appeared as a line of spray on the ground, parallel to the crop row. Hence, spray treatment of small and large weeds could be visually discerned.

The 150025 nozzle tips deposited substantially less dye-water mixture that was difficult to see further away, even after adding more dye to the water mixture. However, the nozzles maintained good coverage when viewed up close. In practice, nozzle tips will be selected and positioned to cover the width of the crop row.

Figure 6.5: The nozzle tips used: TeeJet 1504 (*left*) and UniJet 150025 (*right*)

Methodology for spray application evaluation

The tractor towing the 1.8 m multi-row prototype was driven at 0.8 m/s (3 km/h) over a total of 242 pyrethrum plant to evaluate the prototype's capability to detect and spray weeds in the crop row. The number of weeds hit and pyrethrum plants hit were recorded (true and false positives, respectively). A hit occurred when a plant (i.e. weed or pyrethrum) was in line with the delivered sprays (Figure 6.6).

Figure 6.6: A hit with blue dye with the arrow showing the direction of the row and rectangle showing the kill zone

Results of spray application evaluation

The table below shows the results from the evaluation. Of the 242 pyrethrum plants, 30 were sprayed incorrectly (false positives). Fifteen of the false positives were edges and sections of pyrethrum plants (Figure 6.7), which were not expected to result in a lethal dose to the plant. The other 15 false positives were of pyrethrum that had delayed regrowth (Figure 6.8). False positives were observed to be due to large colour differences and un-established canopy during regrowth. Incorporating data of pyrethrum plants with abnormal appearance into the training set will minimise the number of false positives.

A total of 135 weeds were correctly sprayed. Many of the weeds were hit with high precision (Figure 6.9). The overall weed hit rate was evaluated to be 75% and is considered an exceptional result because:

- Weeds varied greatly in appearance, ranging in size from less than 1 cm to greater than 30 cm in diameter, and often overlapped with the pyrethrum plant edges and one another. Further analysis will determine the distribution of detection rate versus weed size.

- 5-10% of misses were due to the spray solenoids rapidly switching and not turning on in time (i.e. the weed was detected but not precisely sprayed – this will be addressed by fine tuning the spray solenoid parameters).
- 5% of misses were too small to visually discriminate as pyrethrum or weed.

Table 6.1: Detection results

Detection result	Count	Percentage
Total pyrethrum plants	242	
Pyrethrum hit (delayed regrowth)	15	6%
Pyrethrum hit (tips and plant)	15	6%
Total pyrethrum not hit	212	88%
Total pyrethrum hit*	30	12%
Total weeds hit	135	75%**

** Pyrethrum that varied greatly in appearance and that had its tips (i.e. edges) sprayed are included in the false positives*
*** Due to variability in weed size from <1 cm to greater than 30 cm this value is an estimate based from video data recorded during the evaluation.*

Figure 6.7: False positives: edges and small sections of pyrethrum plants.

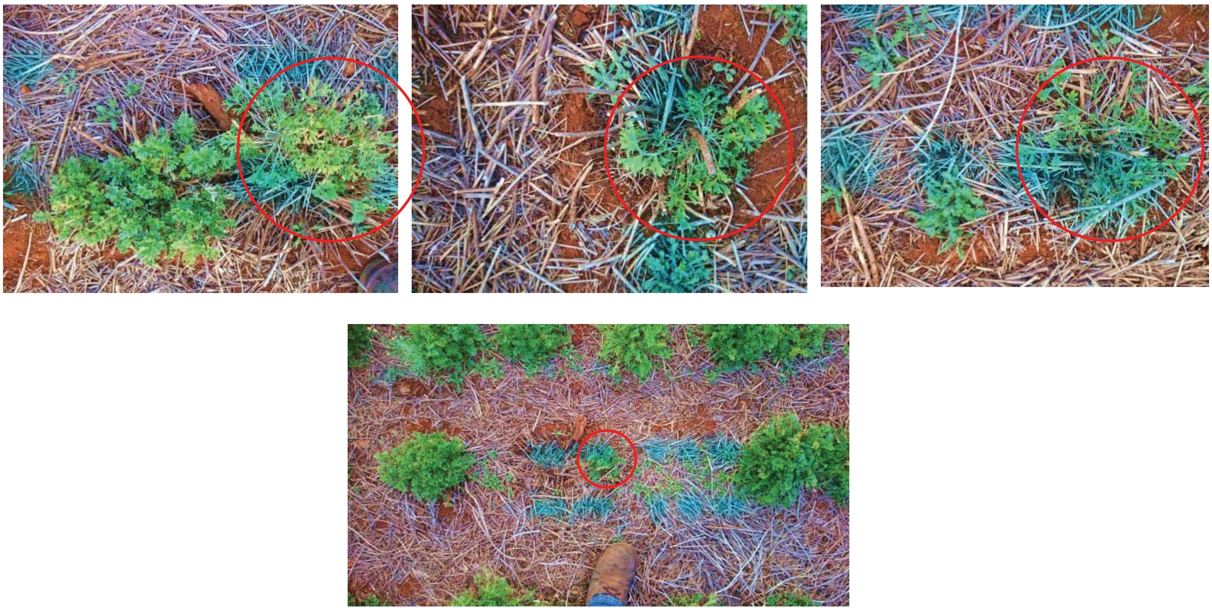

Figure 6.8: False positives: pyrethrum plants regrowth.

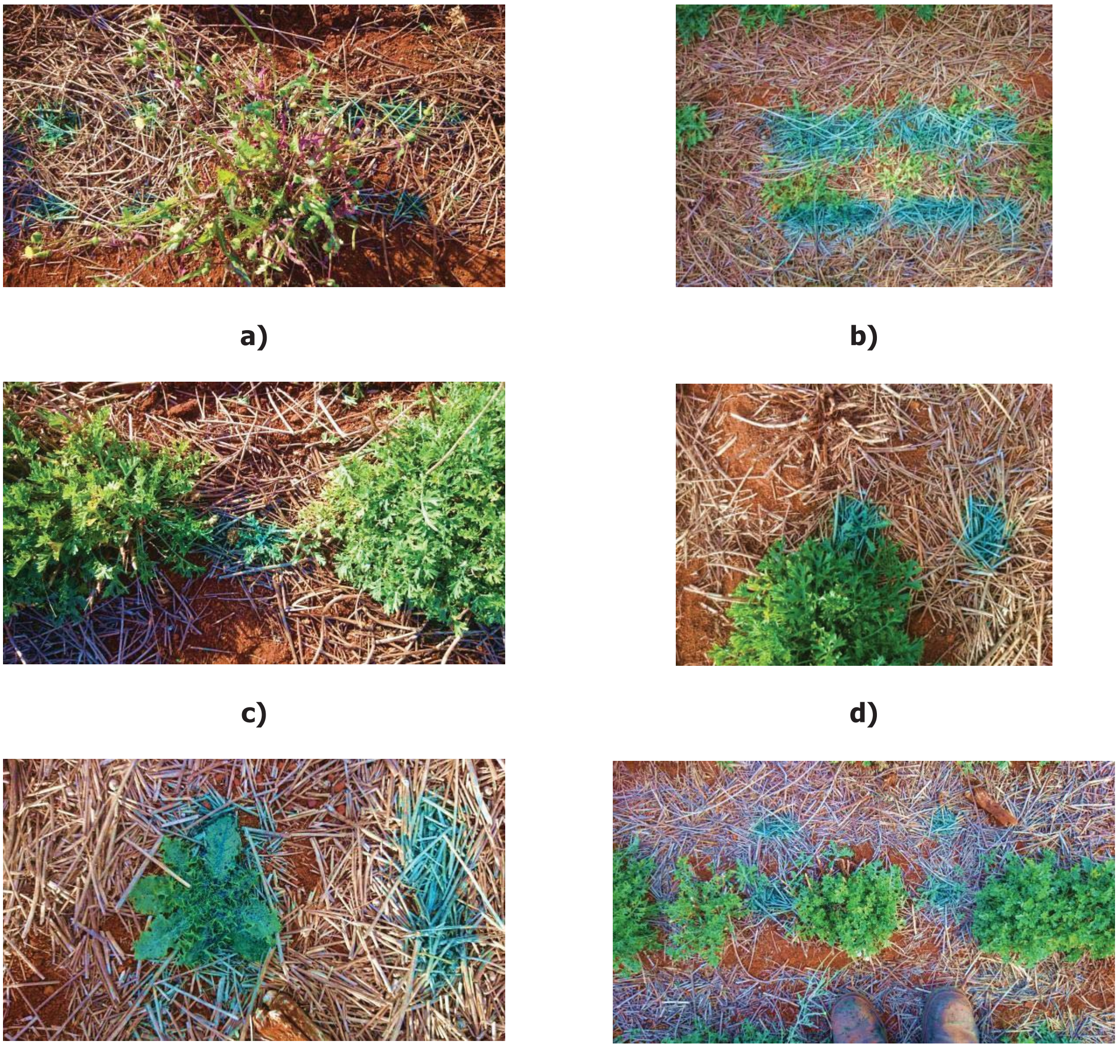

a)

b)

c)

d)

e)

f)

Figure 6.9: True positives: a) large weed ; b) small groundsel; c) small groundsel between pyrethrum plants; d) small groundsel against pyrethrum plant; e) broadleaf weed; f) groundsel between pyrethrum plants

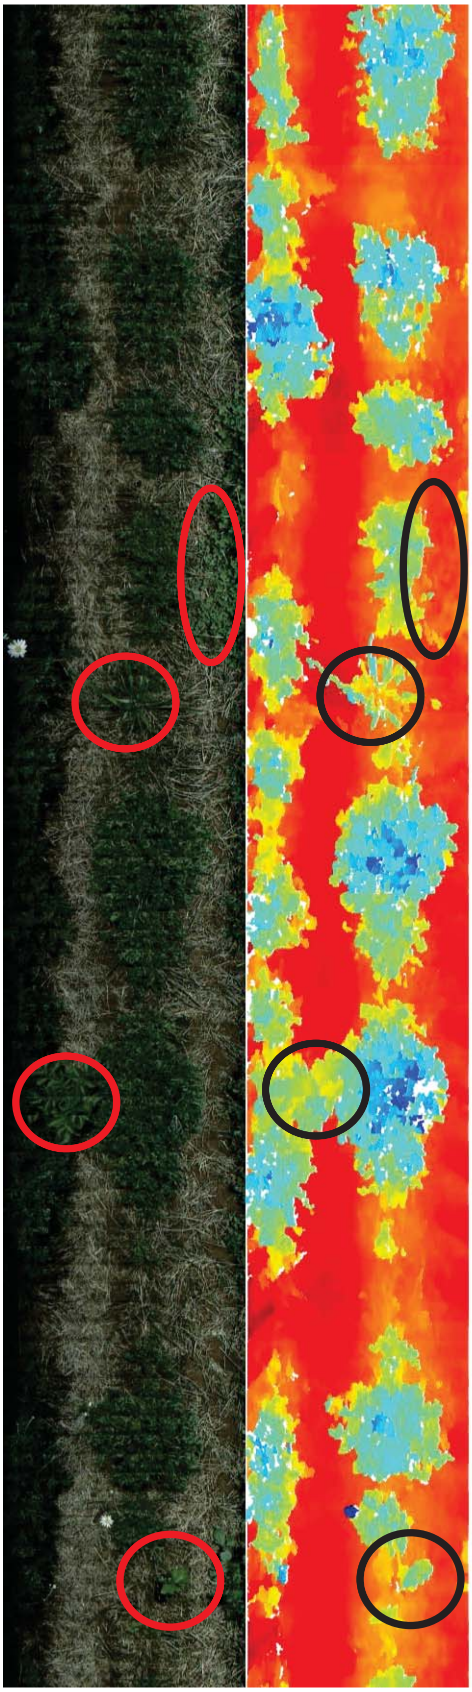

Section of weed map illustrating how weed type (i.e. grass, potato, clover) weeds can be reviewed and identified

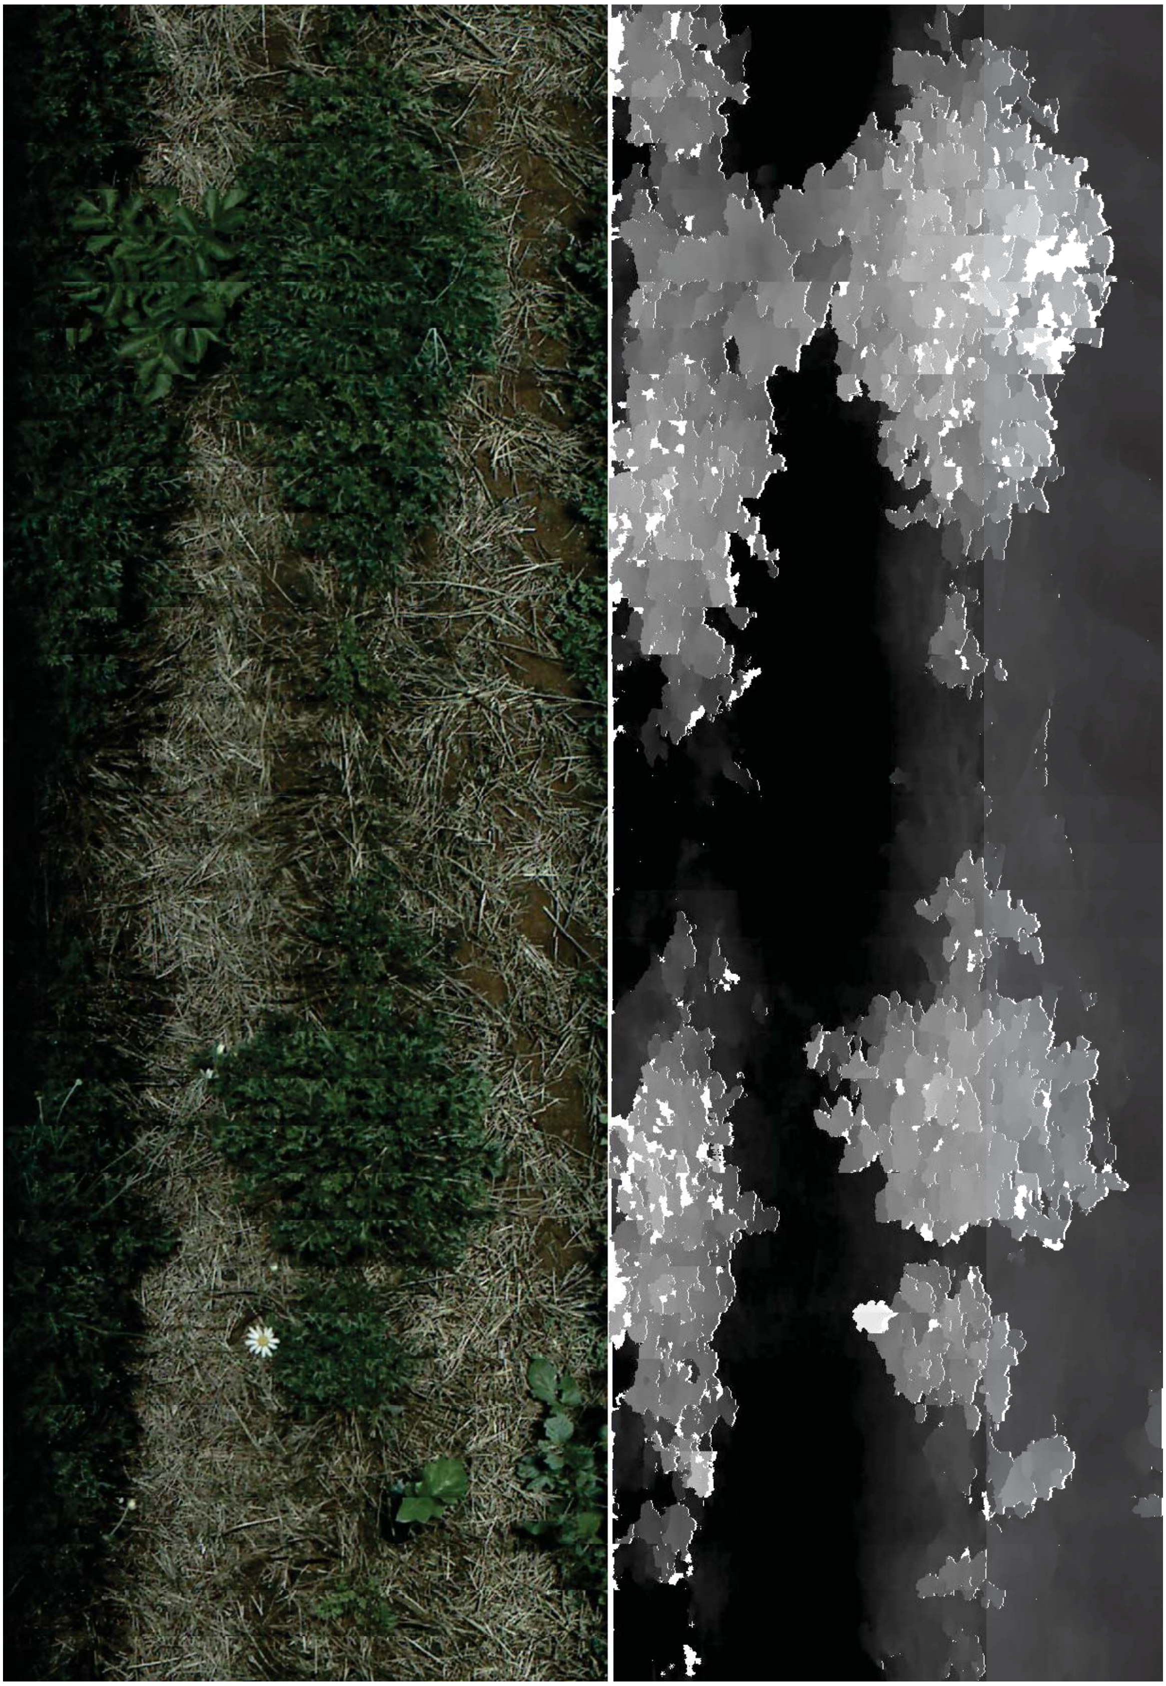

Weeds growing on-the-row and weeds growing between-the-row, weeds butted up against pyrethrum, some of the pyrethrum was in flower

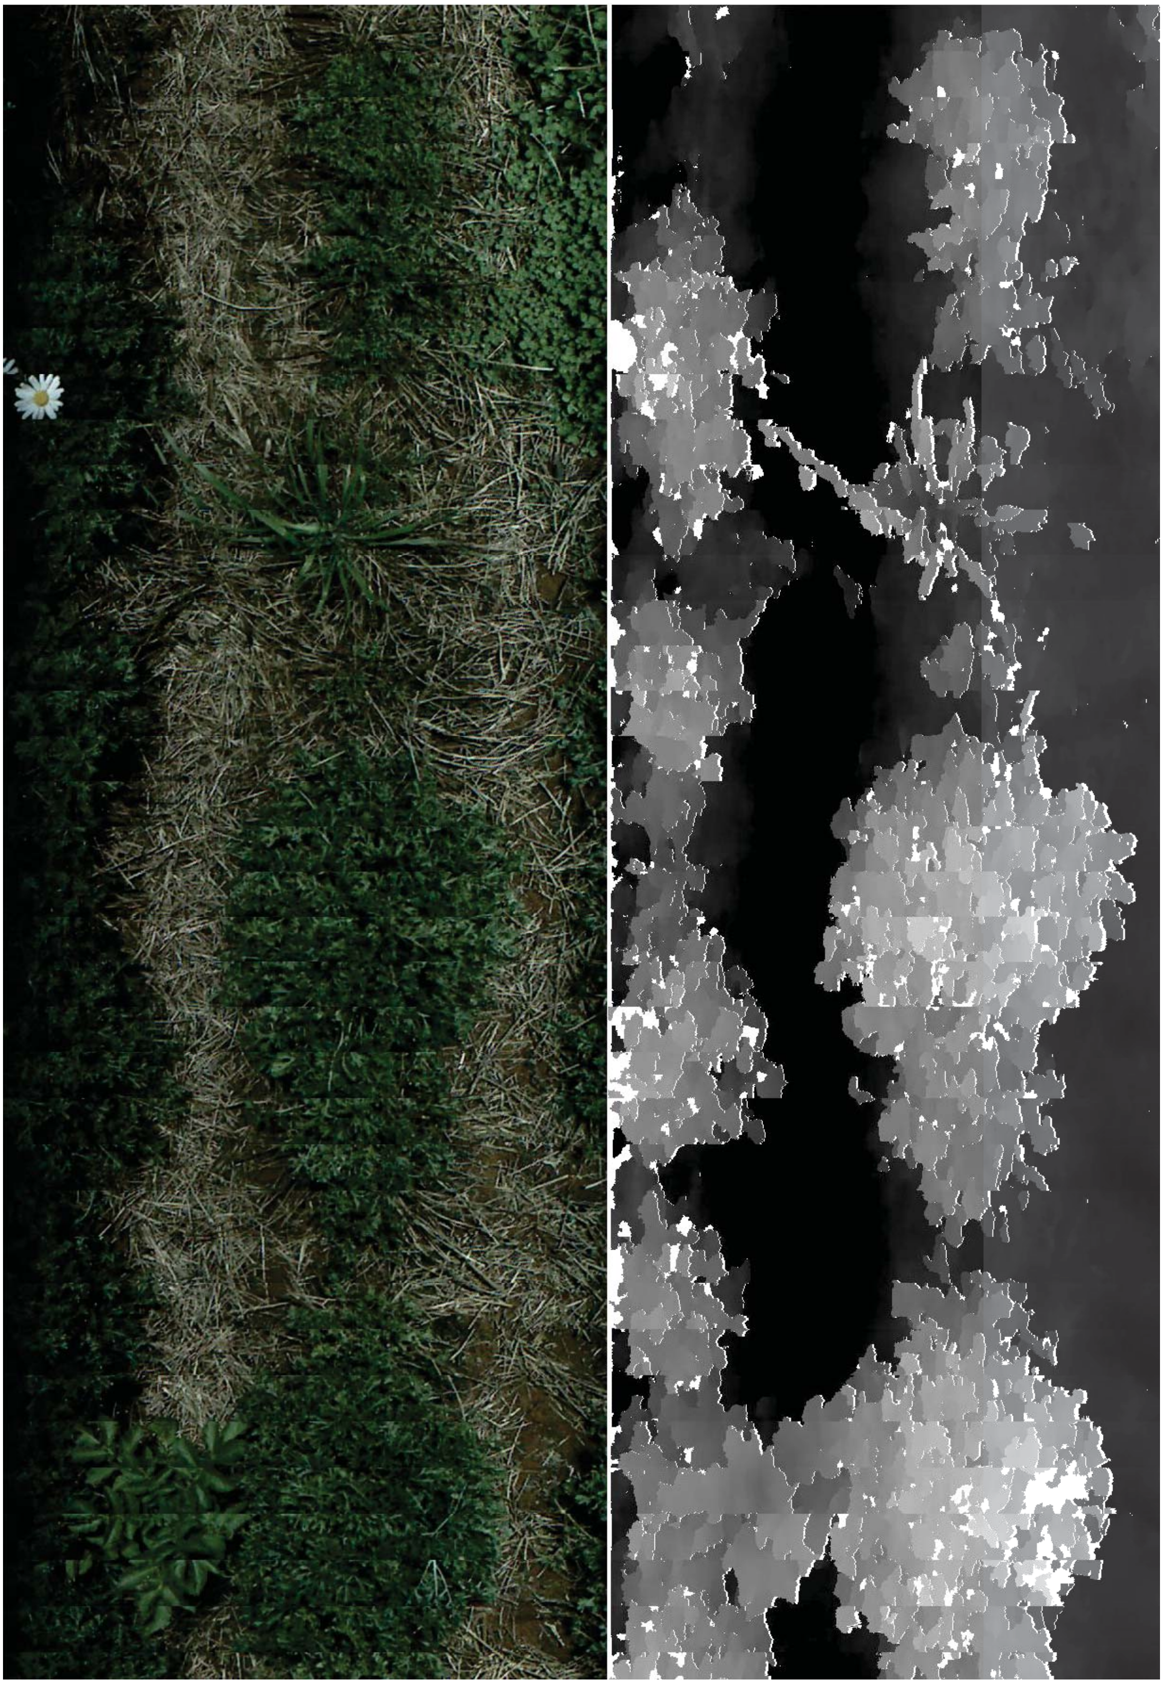

Variety of weeds on and between the row

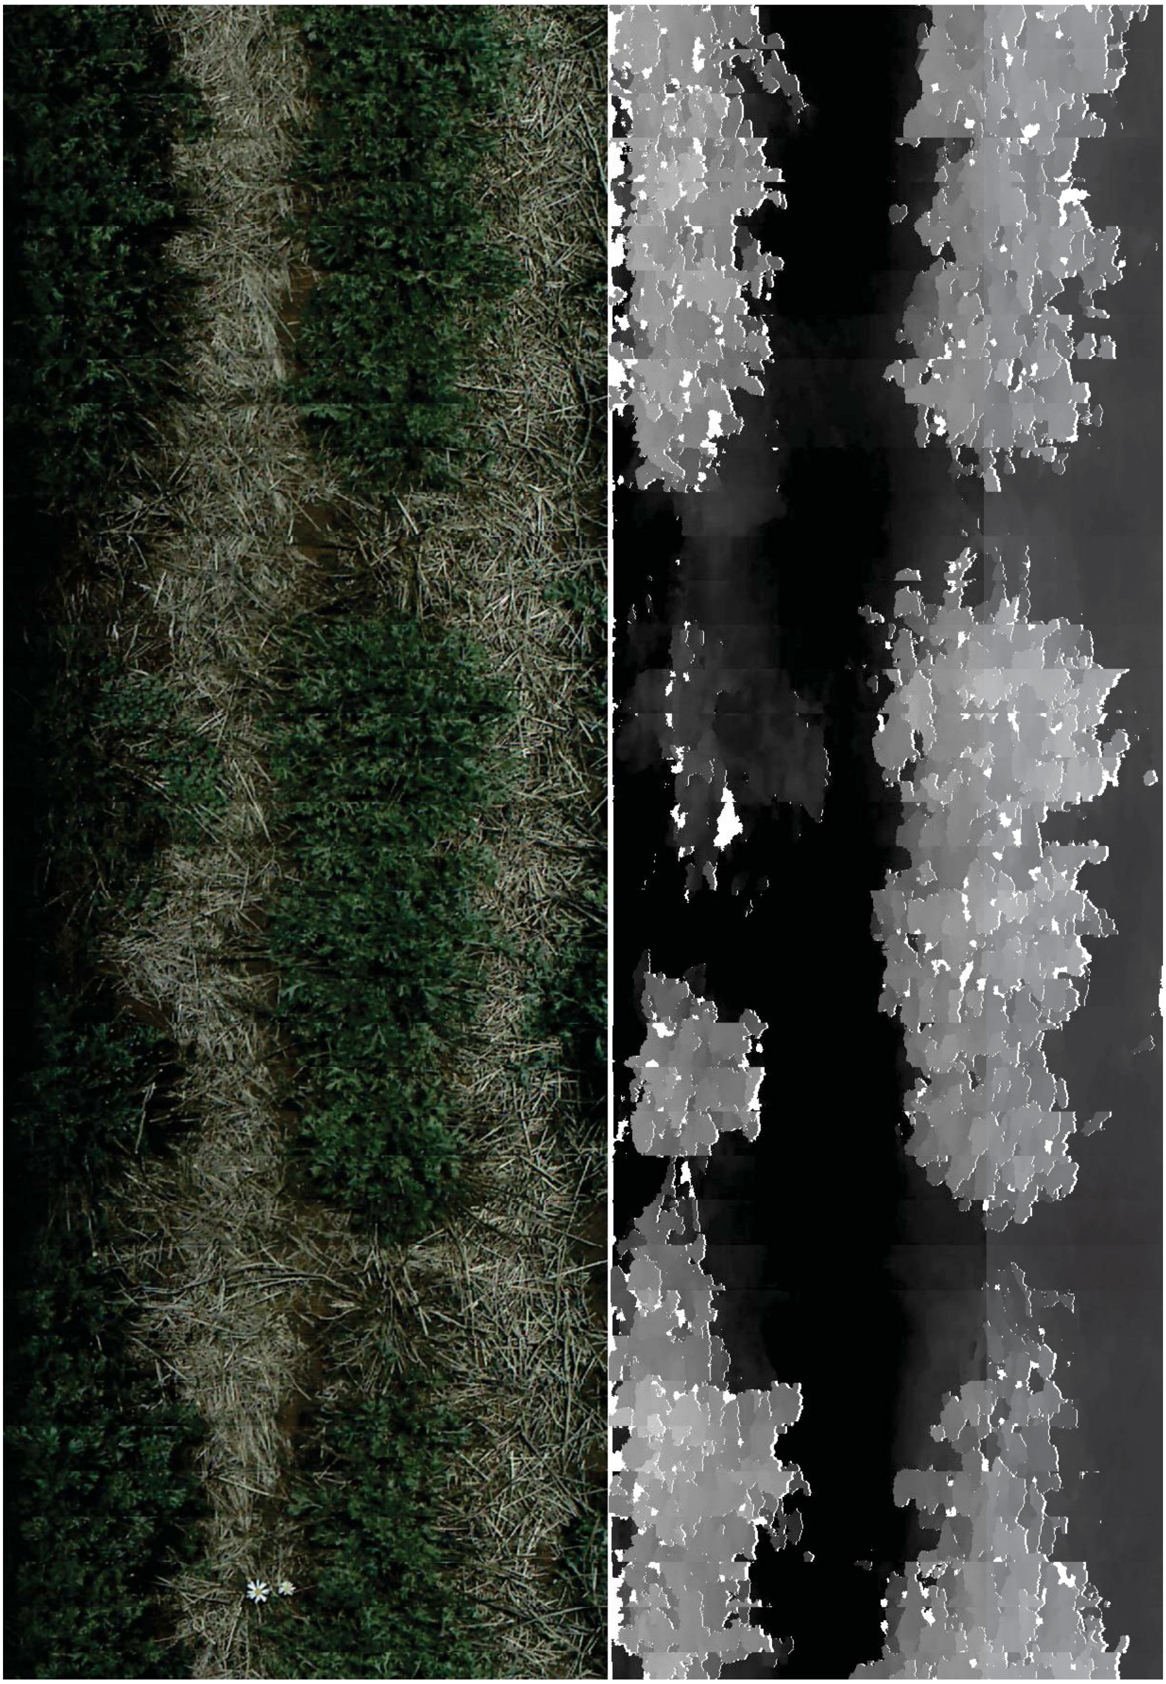

Weeds growing amongst the pyrethrum

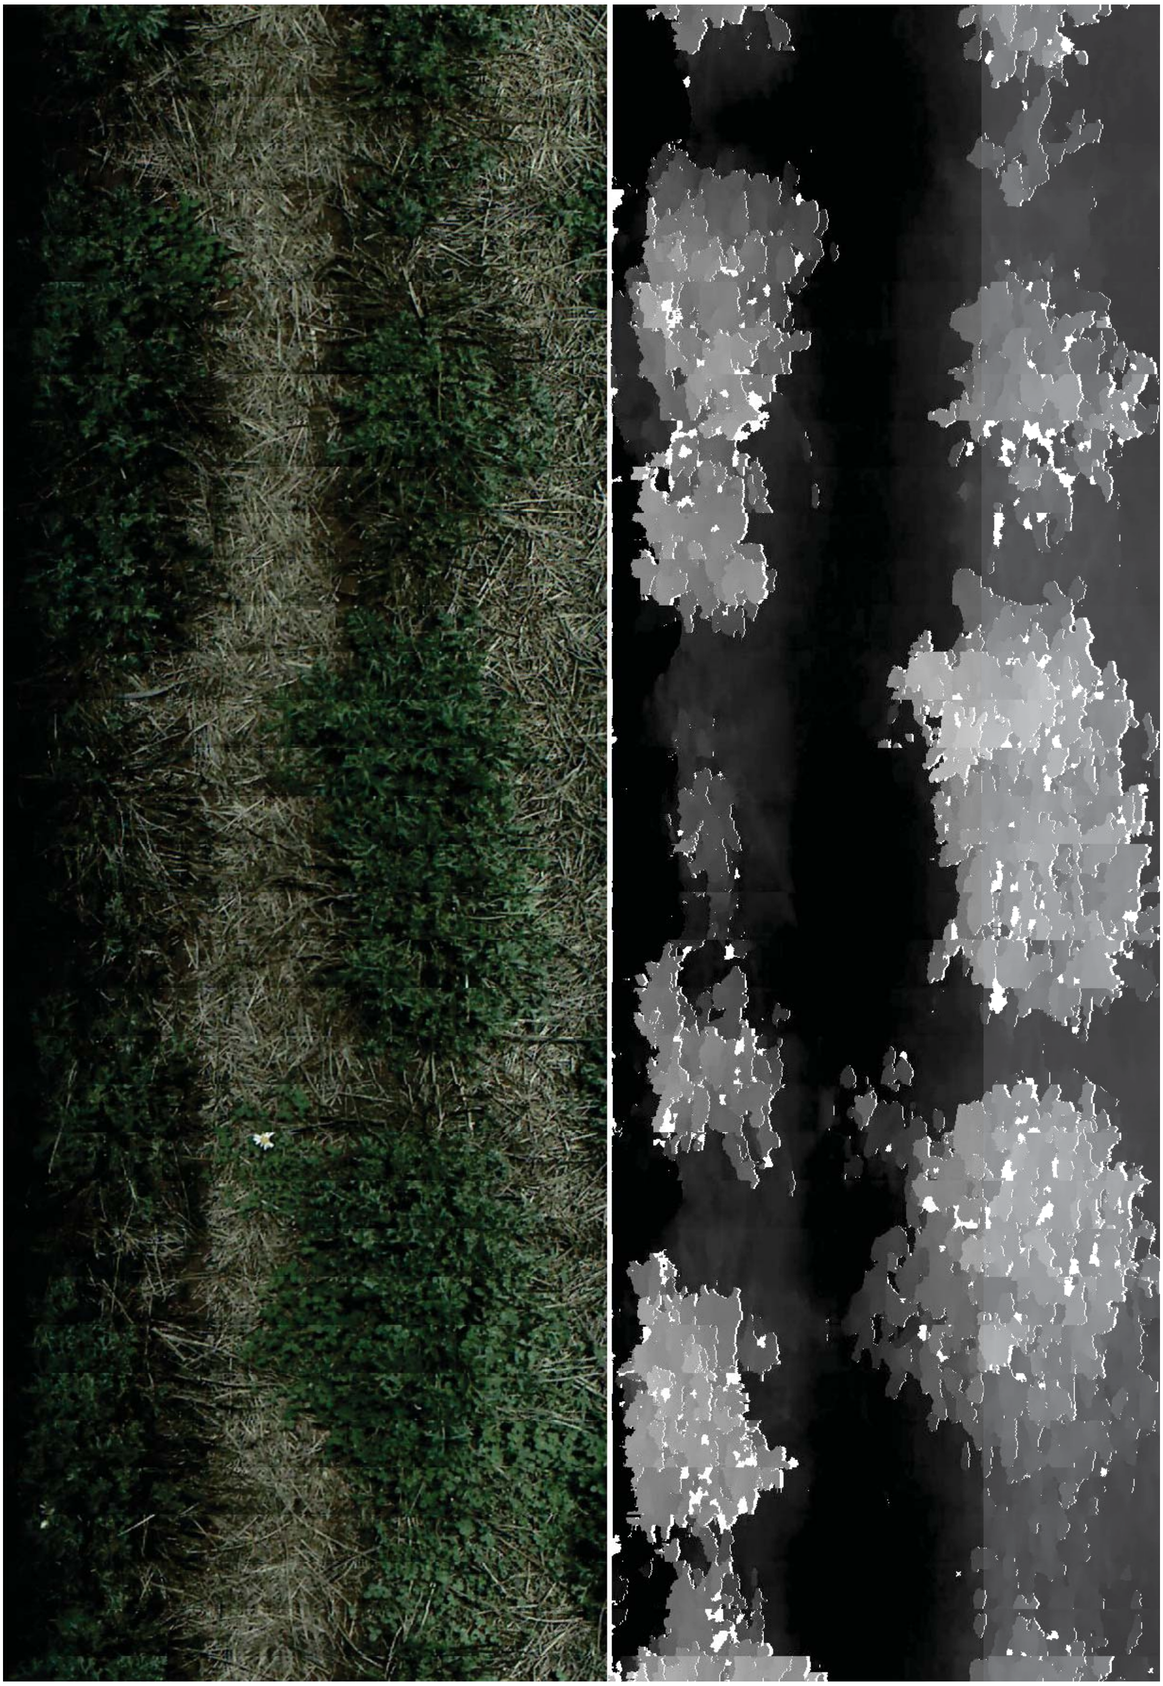

Some of the pyrethrum was so close to the tractor tread that it fell into the wheel track this caused the depth and colour of the pyrethrum to be distorted

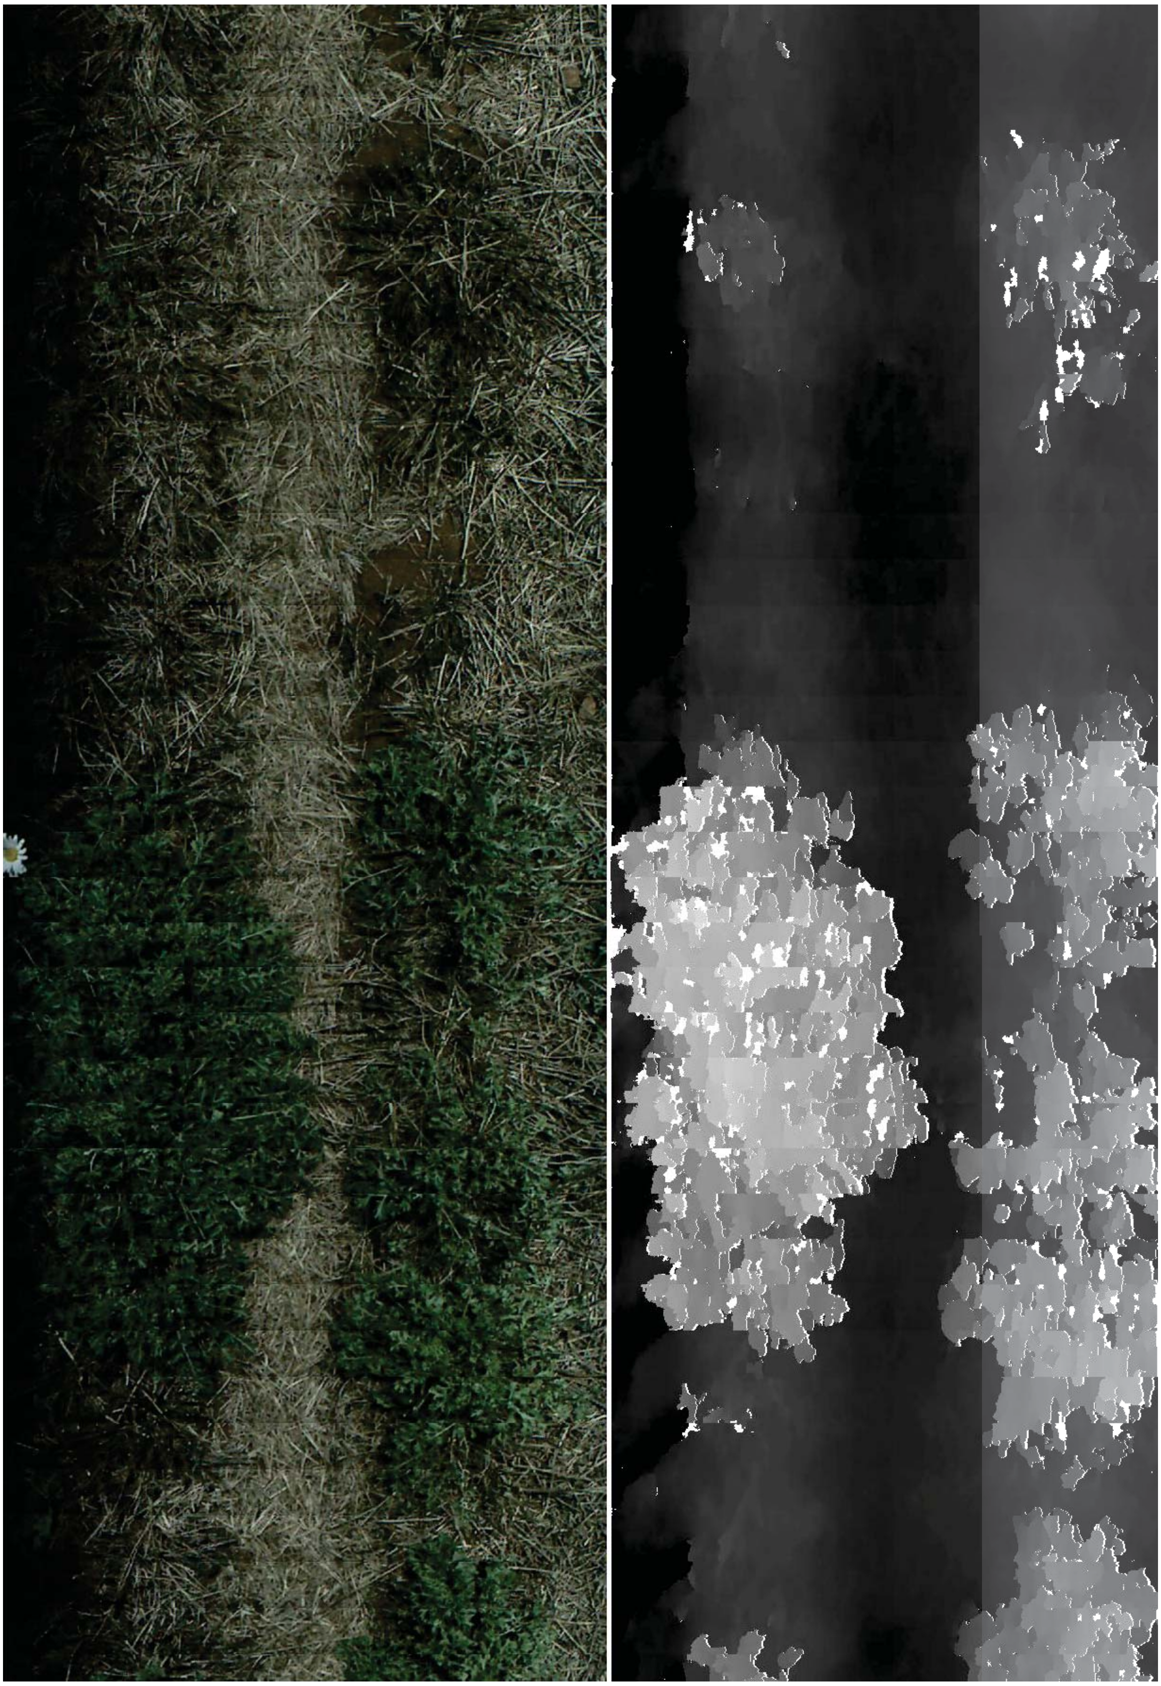

Pyrethrum re-grew at different rates, pyrethrum plant size and shape varied considerably

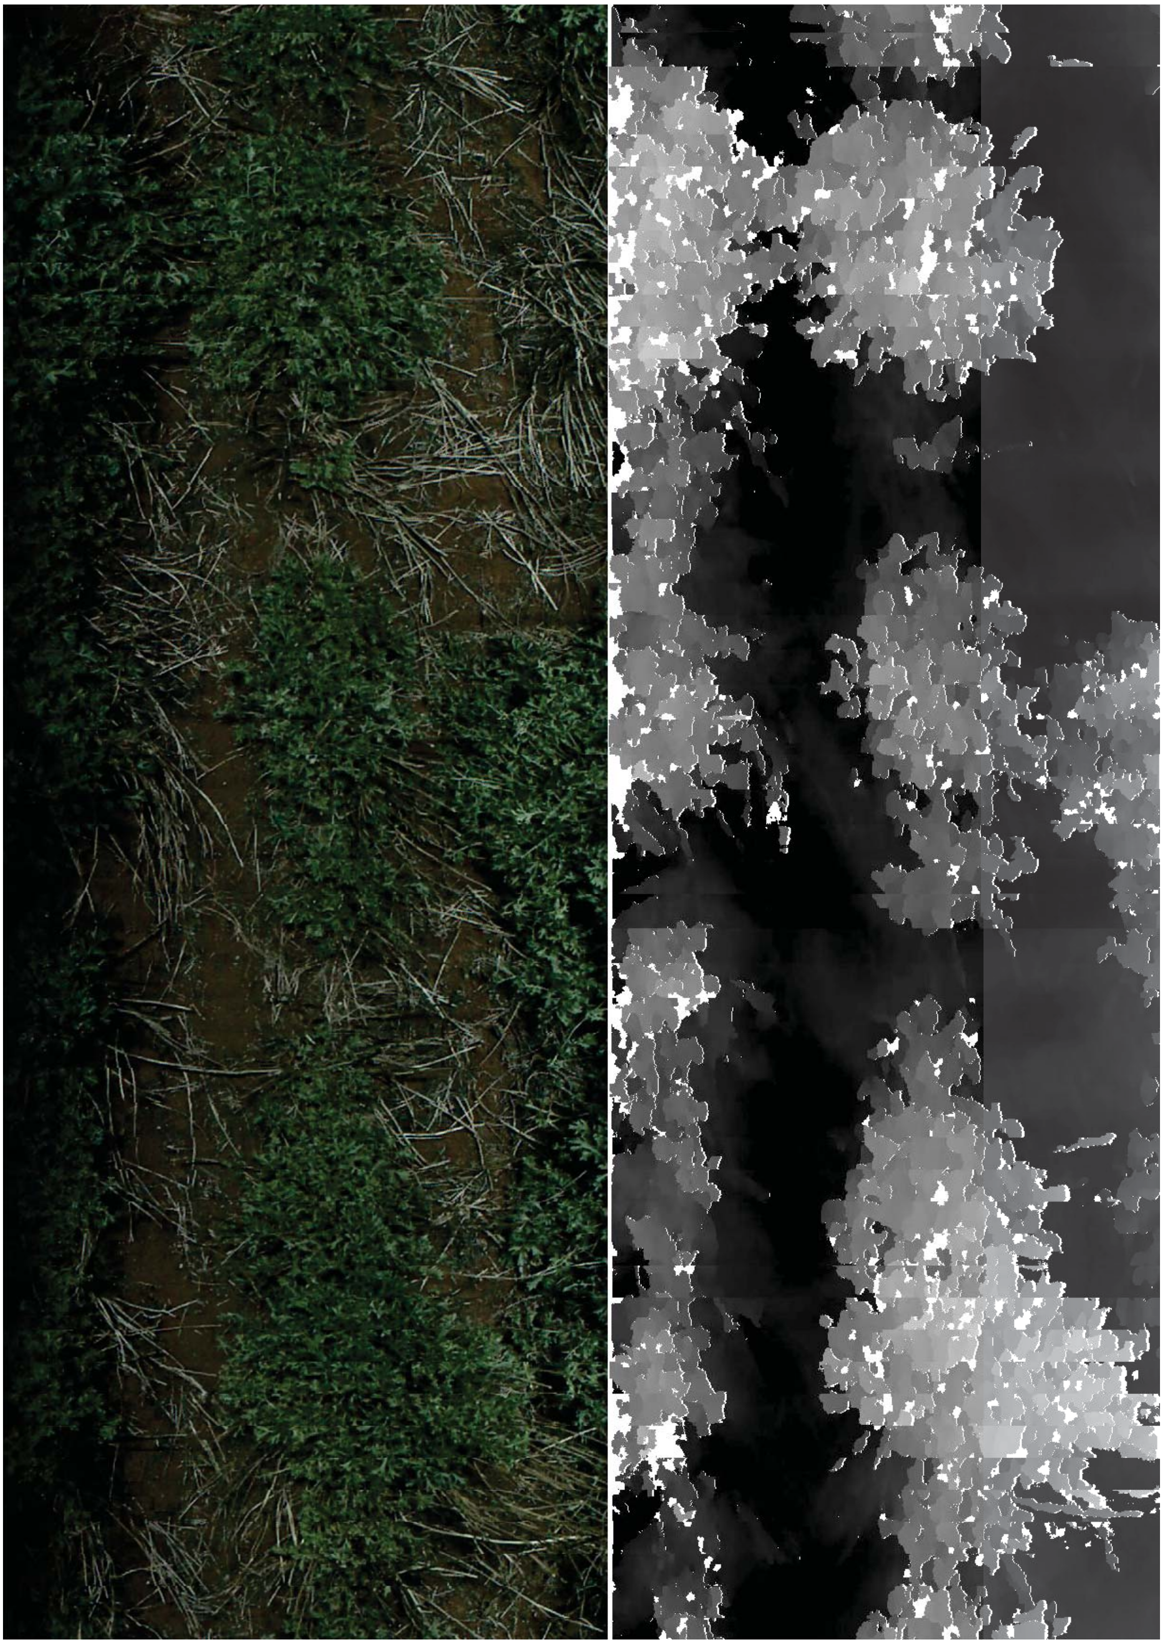

Stubble and soil background

Appendix B – Usage guidelines for precision weed sensing system and spot sprayer

Appendix C – Thesis extracts

Usage guidelines

The ideal precision weed sensing situation was when the pyrethrum crop had:

- ✓ consistent green colour of pyrethrum canopy
- ✓ consistent width, height and shape of pyrethrum
- ✓ consistent row spacing

Non-ideal situations might affect performance of the precision weed sensing. However, the precision weed sensing system catered for a wide range of pyrethrum and weed conditions, as listed below.

1. Soil or stubble background.

The pyrethrum crop had a soil background before harvest, and straw-like pyrethrum stubble background after harvest. The precision weed sensing caters for both backgrounds.

Pre-harvest and post harvest ground conditions

2. Detection of on-the-row and between-row weeds.

Weeds growing in isolation on-the-row and between-row were straightforward to detect.

Weeds growing on-the-row and between-the-row (circled)

3. Consistent and straight crop rows.

The guidance system will perform best in consistent and straight crop rows. The guidance system can identify and ignore short runs of discrepancies to maintain the correct row trajectory, but may experience degradation in performance for long runs of inconsistent rows.

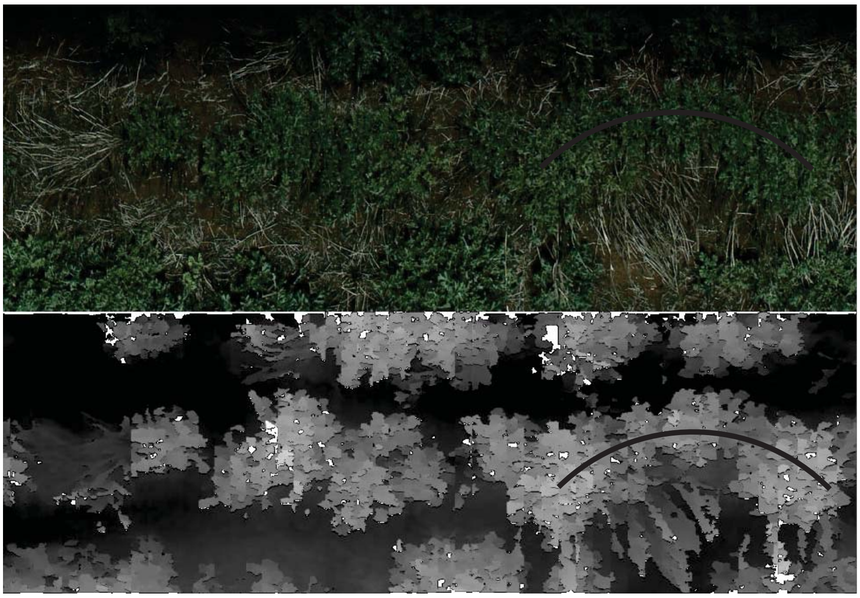

Bend in the row marked in a black line

4. Pyrethrum plants in wheel tracks.

Plants on wheel track rows were observed to have distorted depth, texture and colour from tractor damage or from growing into the wheel track. Rows on the wheel tracks were excluded from analysis.

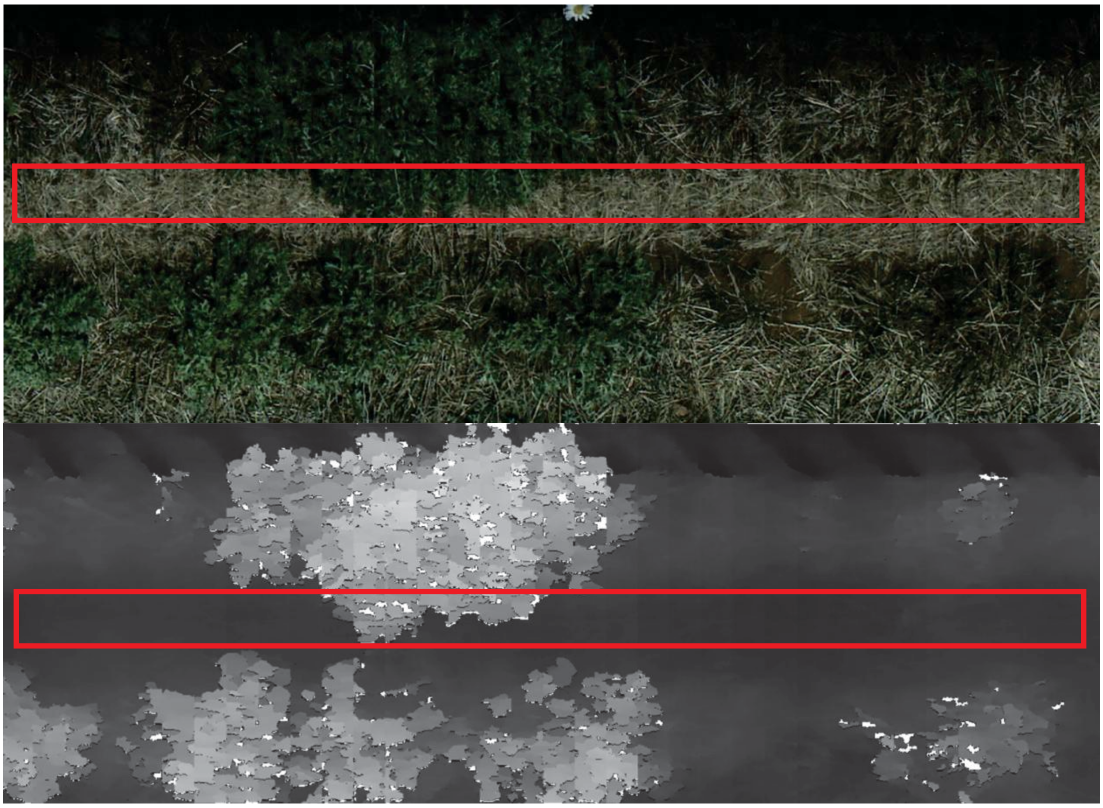

Damaged pyrethrum in the colour image and wheel track in the depth image (red rectangles)

5. Re-growth amongst variable size of pyrethrum stubble.

Variable size of stubble was caused by the harvester cutting off pyrethrum daisies at variable heights due to undulating terrain. Some beds had stubble in excess of 20cm high whilst adjacent rows had stubble 5cm high. The precision weed sensing system catered for variable amounts of stubble, but performance degraded with increasing amounts of stubble. The stubble was grey-brown in colour which fragmented the appearance of green pyrethrum plants.

Stubble fragmented the appearance of pyrethrum re-growth ▶

Variable amount of pyrethrum stubble as demarcated by the black lines

6. Discolouration of pyrethrum tips.

Similar to re-growth amongst stubble, discolouration of pyrethrum tips caused the appearance of the green pyrethrum plants to be fragmented. Discolouration of pyrethrum tips was caused by:

- grey-brown tips from dieback after broadcast spraying and the onset of winter
- yellow tips from other plant conditions

Pyrethrum with grey-brown tips and yellow tips

7. Early flowering pyrethrum plants.

A large height contrast occurred when some of the pyrethrum plants flowered early and had tall daisies and foliage. Non-flowering pyrethrum plants were significantly lower.

Large amount of tall daisies and foliage, which were several centimetres above non-flowering pyrethrum plants, generated false positive triggers intended to target taller weeds.

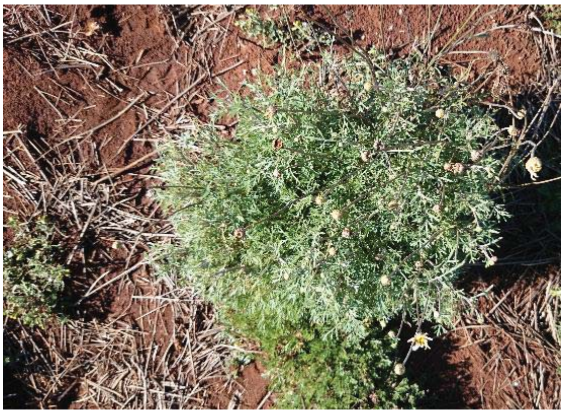

Flowering pyrethrum plant (centre) next to a smaller non-flowering plant (bottom centre)

8. Variable pyrethrum size and colour.

While some plants were yet to re-sprout after harvest, others were flowering prematurely, leading to substantial height, color, texture, size and shape variation. The precision weed sensing system compensated for variations in pyrethrum appearance.

Pyrethrum ranging dramatically after regrowth (circled) in a groundsel infestation

Pyrethrum plants demonstrated variable colour from light green, green and blue green

9. Diverse weed spectrum and appearance.

The precision weed sensing system discriminated between pyrethrum and weeds by classifying a plant as pyrethrum or not pyrethrum. This is because the weed spectrum was diverse and included weeds that were:

- sparse, e.g. grasses
- high and dense, e.g. potatoes and mallow
- low and dense, e.g. clover
- textured, e.g. bur-chervil
- of a different colour but similar texture or shape to pyrethrum, e.g. pink weed

White clover infestation with an affected pyrethrum plant (circled)

Large grass weed on the row

Bur-chervil (lighter green) amongst pyrethrum canopy

Broadleaf weed butted up against a pyrethrum plant on-the-row (circled)

10. Sunlight effects on image lighting.

The hilly terrain caused variable amounts of sunlight to infiltrate under the skirting used to control the lighting for certain times of the day and implement orientations relative to the sun. This lighting variability caused shifts in detected green colour. The skirting on the prototype was modified for the purpose of blocking out and removing the effects of sunlight on the precision weed sensing system.

Table 3.1: Sugarcane data collection at 'Fairymead', Bundaberg.

Date	**Location**	**Run length (m)**	**Crop height (m)**	**Variety**	**Growth stage**	**Ratoon**	**Trash blanket**	**Speed (km/h)**
19/6/2012	27-A	70	0.5-0.8	Q208	medium	3	yes	2.5
5/9/2012	13-B	100	0.1	Q151	short	2	no	2.5
10/10/2012	28-B	991	0.25	Q151	short	2	yes	2.5
10/10/2012	13-A	1,747	0.8-1.0	Q232	medium	3	no	3.5
6/11/2012	2-A	2,260	1.3	Q151	high	3	no	5
6/11/2012	4-B	2,643	0.25-0.5	Q208	short/ medium	4	yes	5
4/12/2012	27-A	12,213	1.0-1.3	Q208	medium/ high	3	yes	5

3.5 Pyrethrum data acquisition system

A single crop-row data acquisition system (1 m wide × 1.6 m long × 1 m high) shown in Figure 3.9 was built and instrumented with cameras, lights and computer equipment set out in Section 3.2 and shown in Figure 3.10. For data collection in pyrethrum, the data acquisition system included a light-restricting cover to allow the data acquisition system to operate during the day. The camera and lights were mounted at a height of 1 m above the ground which provided a Region Of Interest (ROI) on the ground of 1 m wide and 0.6 m long. A field data collection program incorporating pyrethrum's growth cycle, weed infestation and different growing conditions was developed (Section 3.5.1) and the single crop-row data acquisition system was deployed to Botanical Resources Australia[3] (BRA) who facilitated the data collection program.

---

[3] http://www.botanicalra.com.au/

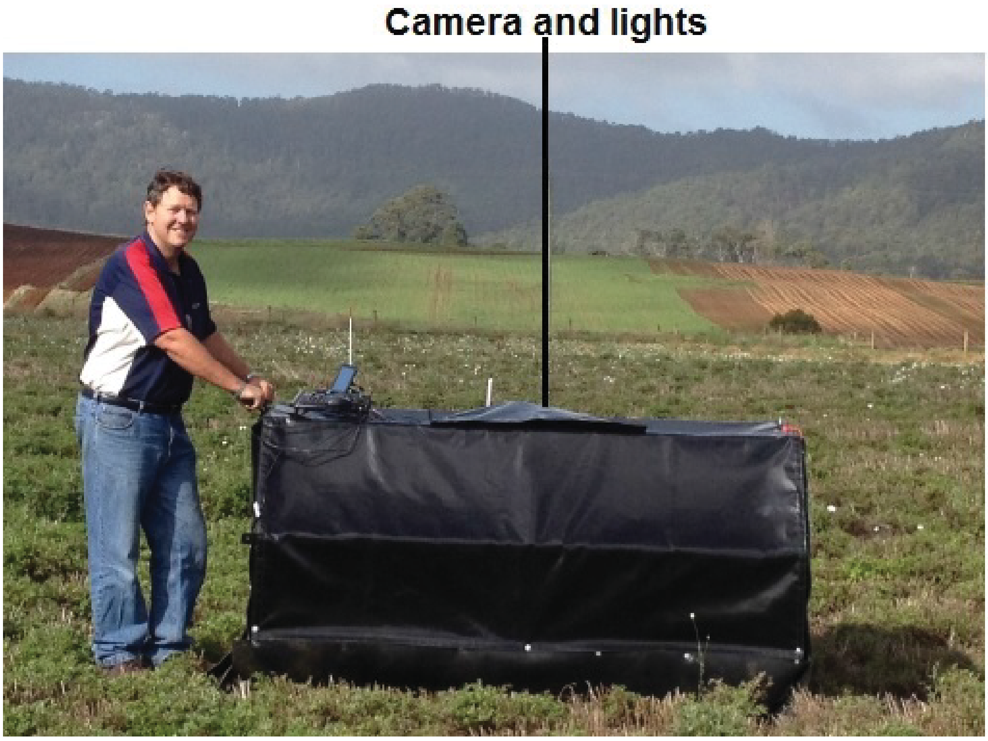

Figure 3.9: Pyrethrum data acquisition system (1 m wide × 1.6 m long × 1 m high) in a crop of pyrethrum March 2012. Author (1.86 m tall) to indicate scale.

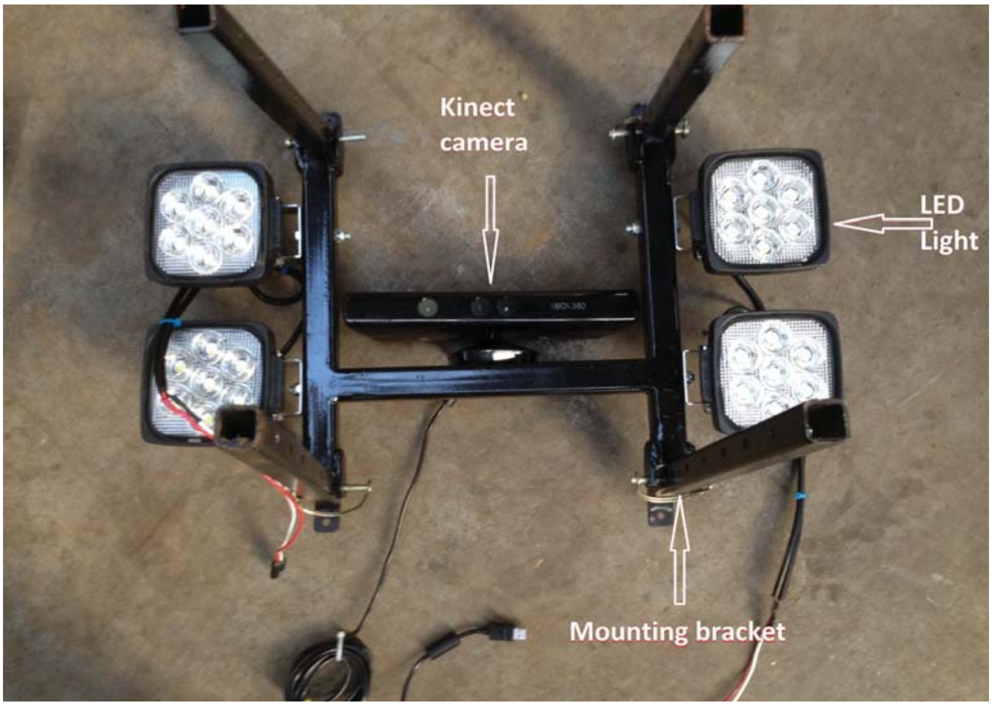

Figure 3.10: Kinect$^{\circledR}$ camera (centre) and lights (four off) for mounting on the pyrethrum data acquisition system. View from below.

3.5.1 Pyrethrum site criteria and field data collection

Site selection

The ideal selection of sites are those that cover the dominant soil types in the pyrethrum industry of northern Tasmania (i.e. black, brown and sandy). Each soil type should contain an area of differing stubble cover (i.e. trash and no trash) and the typical weed spectrum. Pyrethrum is the target plant for identification but inclusion of the full weed spectrum in the data set is expected to increase robustness of pyrethrum and non-pyrethrum classification.

Collection timing

The ideal interval between data acquisition events would capture the pyrethrum and weed data at all differing growth stages during the typical weed control period. In consultation with the BRA research and development group, the data acquisition interval was determined to be seven days. If a data acquisition event was missed, because of rain or other weather events, then the earliest time the operator could access the field sites after the weather event would suffice.

Data collection groundspeed

Herbicide spray application occurs at groundspeeds from 3 km/h up to 8 km/h for pyrethrum. Therefore the data acquisition system groundspeed should be comparable to provide real-world evaluation.

3.5.2 Agronomic factors affecting pyrethrum data

3.5.2.1 Die-back

Data collected in 2012 highlighted the presence of a disease that caused a 'die-back' in the plant, where parts of the plant die. The disease was an industry-wide problem and is found to be prevalent in wetter years. Figure 3.11 presents an image of a healthy pyrethrum plant compared to a diseased pyrethrum plant,

Figure 3.12. BRA began researching the agronomic implications and control of the die-back, as the die-back created problems in both the agronomic system (e.g. poor growth vigor) and the machine vision systems algorithm development (e.g a single plant component was split into several smaller components). In consultation with BRA, it was decided that the weed control strategies should be restricted to a time frame between April and early June (the post harvest vegetative growth period), as this is when the die-back was least noticeable.

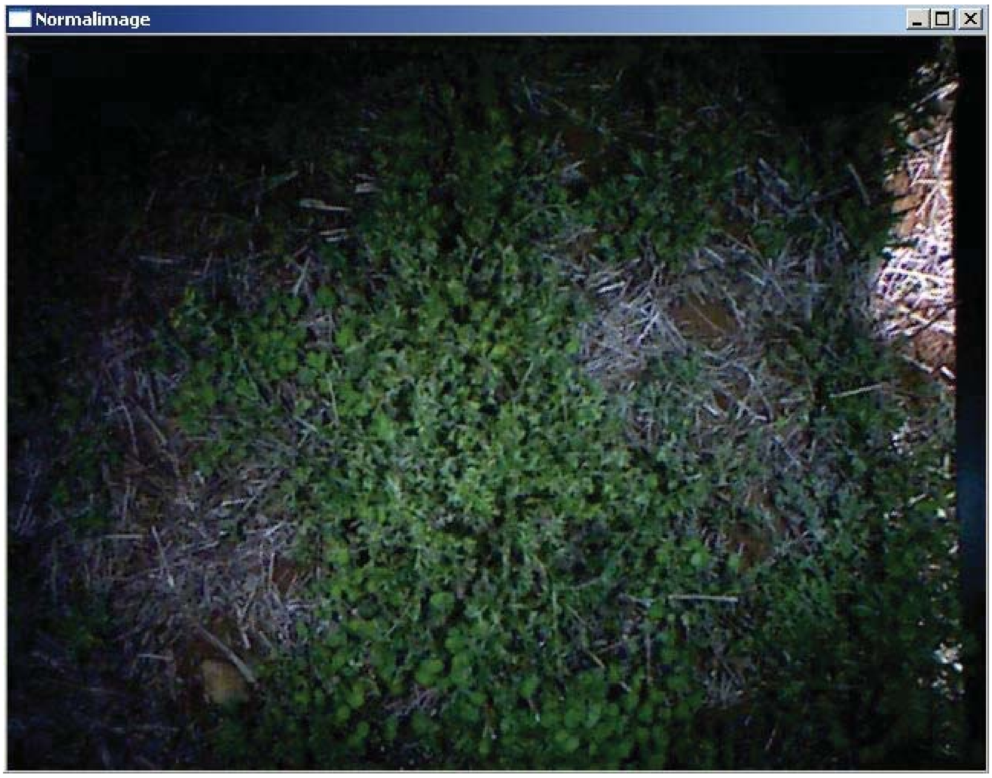

Figure 3.11: Healthy pyrethrum plant at the post-harvest semi-dormant growth stage.

Figure 3.12: Unhealthy pyrethrum plant exhibiting ‘die-back’ at the post-harvest dormant growth stage.

3.5.2.2 Row spacing effect on plant size

Commercial spot spray herbicide applicators (solenoid nozzles) are incapable of spraying weeds without causing overspray onto the crop at low crop row spacings (rows planted less than 0.25 m apart). Therefore BRA planted trial plots of pyrethrum at a range of row spacings from 0.4 m to 0.7 m. The data acquisition system was used to obtain data from each of the row widths. Cole’s, Dick’s and BRA Jamison’s sites were planted with a row spacing of 0.2 m and Gibson’s site was planted on 0.2 m, 0.3 m and 0.4 m rows for yield trials. All collected data was used for observations of visual attributes of weeds to aid in algorithm development, but was not applicable for testing of algorithms based on the relative spatial positioning of the crop and weeds in wider row spacing.

The DRF speedlings site was planted with 0.65 m row spacing and used for testing and trialling of the weed detection algorithms. Pyrethrum plants were observed to grow larger (width and height) and have more variable plant-to-plant size, in the wider row spacing than the narrow row spacing, which was confirmed by BRA agronomic research staff. Data collected using sensor (Kinect® camera system) heights of 0.85 m and 1 m revealed that the lower collection height of 0.85 m did not capture the full width of the larger pyrethrum plants for the wider row spacing, hence one metre was the most suitable sensor height

3.6 Pyrethrum data collection

The data was collected from April 2012 to August 2012 and April and May in 2013 with each set of contiguous frames being termed a ‘run’. The data was collected from five sites identified as: BRA Jamison’s, DRF Speedlings[4], Cole’s, Dick’s and Gibson’s. The first four sites were located within a forty kilometre radius of Ulverstone, and Gibson’s was in the Launceston region, all were in Tasmania’s north, which is where the majority of the Australian pyrethrum is grown. Table 3.2 shows the latitude and longitude co-ordinates of the sites. The first four sites were in what BRA considered a higher rainfall area (relative to the annual rainfall in northern Tasmania) with a range of conditions such as soil colour, stubble cover, and slope of the land. Gibson’s site was in a lower rainfall area on a newly planted field with no stubble. Groundspeeds for the data acquisition systems of between 1 km/h and 8 km/h were used. Weeds found across all sites were flat weed, groundsel, thistle, sow thistle, dandelion, white clover, wireweed and red clover. The data collected from each site is displayed in Tables 3.3 to 3.6 with the key for the tables in Table 3.7. Supporting information on the weeds present at each site is set out in Table 3.8.

---

[4] Speedlings are pyrethrum plants that were grown as seedlings and then planted into the field as opposed to seed being directly planted into the field.

Table 3.2: Latitude and longitude position of pyrethrum data collection sites.

Site	**Latitude**	**Longitude**
DRF Speedlings	-41.190812	146.287121
Cole’s	-41.160795	145.995027
Dick’s	-41.173437	146.306808
BRA Jamieson’s	-41.122677	146.085330
Gibson’s	-41.537748	146.900092

Table 3.3: Pyrethrum data collected at the DRF speedling site.

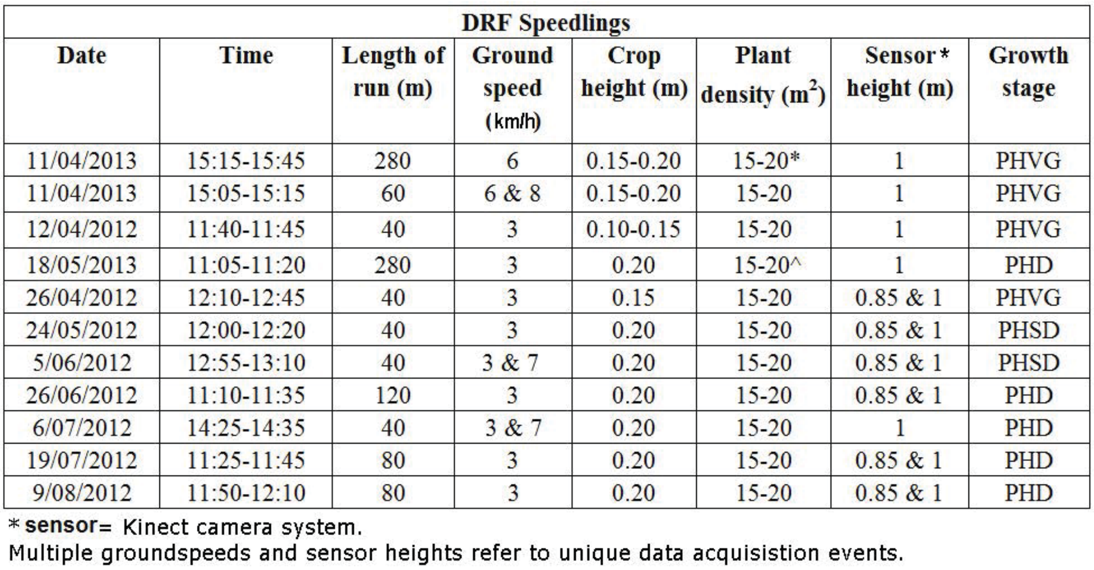

DRF Speedlings

Date	**Time**	**Length of run (m)**	**Ground speed (km/h)**	**Crop height (m)**	**Plant density (m^2)**	**Sensor* height (m)**	**Growth stage**
11/04/2013	15:15-15:45	280	6	0.15-0.20	15-20*	1	PHVG
11/04/2013	15:05-15:15	60	6 & 8	0.15-0.20	15-20	1	PHVG
12/04/2012	11:40-11:45	40	3	0.10-0.15	15-20	1	PHVG
18/05/2013	11:05-11:20	280	3	0.20	15-20^	1	PHD
26/04/2012	12:10-12:45	40	3	0.15	15-20	0.85 & 1	PHVG
24/05/2012	12:00-12:20	40	3	0.20	15-20	0.85 & 1	PHSD
5/06/2012	12:55-13:10	40	3 & 7	0.20	15-20	0.85 & 1	PHSD
26/06/2012	11:10-11:35	120	3	0.20	15-20	0.85 & 1	PHD
6/07/2012	14:25-14:35	40	3 & 7	0.20	15-20	1	PHD
19/07/2012	11:25-11:45	80	3	0.20	15-20	0.85 & 1	PHD
9/08/2012	11:50-12:10	80	3	0.20	15-20	0.85 & 1	PHD

* **sensor**= Kinect camera system.
Multiple groundspeeds and sensor heights refer to unique data acquisistion events.

Table 3.4: Pyrethrum data collected at the Cole’s site.

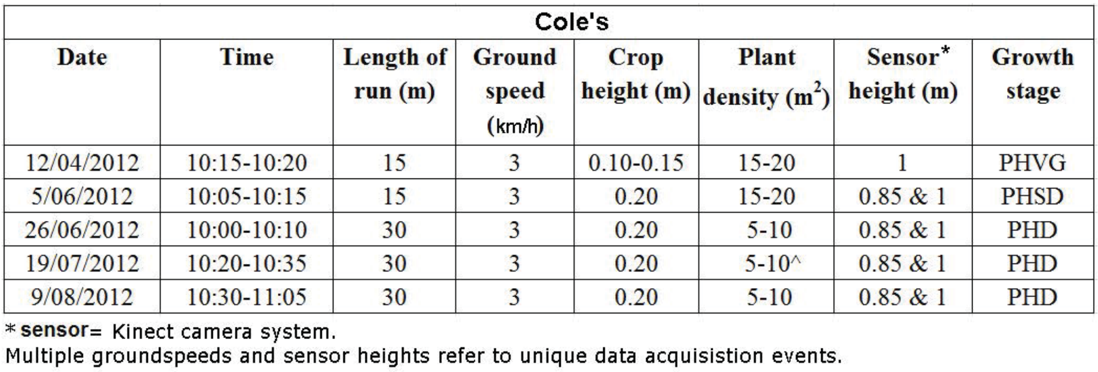

Cole's

Date	**Time**	**Length of run (m)**	**Ground speed (km/h)**	**Crop height (m)**	**Plant density (m^2)**	**Sensor* height (m)**	**Growth stage**
12/04/2012	10:15-10:20	15	3	0.10-0.15	15-20	1	PHVG
5/06/2012	10:05-10:15	15	3	0.20	15-20	0.85 & 1	PHSD
26/06/2012	10:00-10:10	30	3	0.20	5-10	0.85 & 1	PHD
19/07/2012	10:20-10:35	30	3	0.20	5-10^	0.85 & 1	PHD
9/08/2012	10:30-11:05	30	3	0.20	5-10	0.85 & 1	PHD

* **sensor**= Kinect camera system.
Multiple groundspeeds and sensor heights refer to unique data acquisistion events.

Table 3.5: Pyrethrum data collected at the Dick's site.

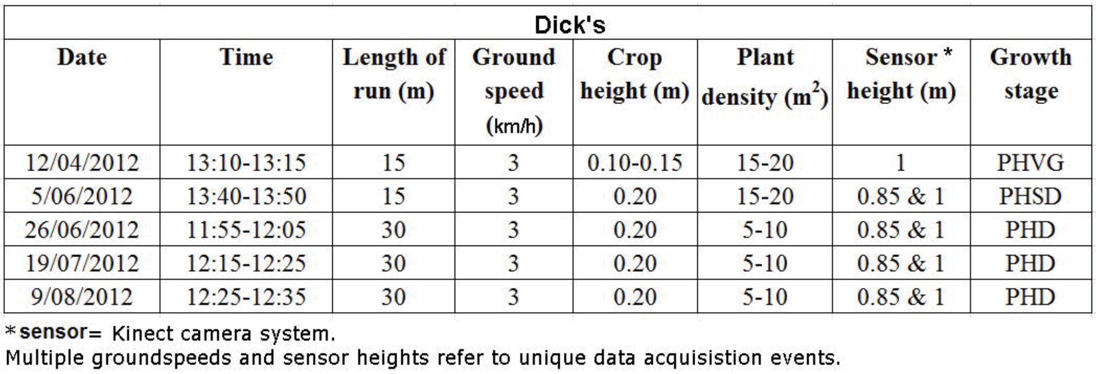

Dick's							
Date	**Time**	**Length of run (m)**	**Ground speed (km/h)**	**Crop height (m)**	**Plant density (m^2)**	**Sensor * height (m)**	**Growth stage**
12/04/2012	13:10-13:15	15	3	0.10-0.15	15-20	1	PHVG
5/06/2012	13:40-13:50	15	3	0.20	15-20	0.85 & 1	PHSD
26/06/2012	11:55-12:05	30	3	0.20	5-10	0.85 & 1	PHD
19/07/2012	12:15-12:25	30	3	0.20	5-10	0.85 & 1	PHD
9/08/2012	12:25-12:35	30	3	0.20	5-10	0.85 & 1	PHD

***sensor**= Kinect camera system.
Multiple groundspeeds and sensor heights refer to unique data acquisistion events.

Table 3.6: Pyrethrum data collected at the Gibson's and BRA Jamison's site.

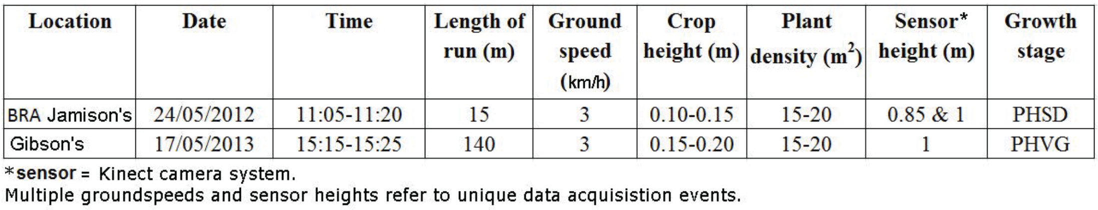

Location	**Date**	**Time**	**Length of run (m)**	**Ground speed (km/h)**	**Crop height (m)**	**Plant density (m^2)**	**Sensor* height (m)**	**Growth stage**
BRA Jamison's	24/05/2012	11:05-11:20	15	3	0.10-0.15	15-20	0.85 & 1	PHSD
Gibson's	17/05/2013	15:15-15:25	140	3	0.15-0.20	15-20	1	PHVG

***sensor** = Kinect camera system.
Multiple groundspeeds and sensor heights refer to unique data acquisistion events.

Table 3.7: Key to accompany Tables 3.3 to 3.6.

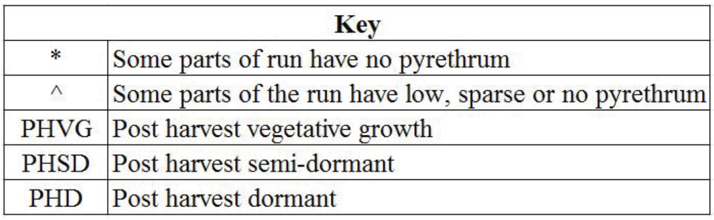

Key	
*	Some parts of run have no pyrethrum
^	Some parts of the run have low, sparse or no pyrethrum
PHVG	Post harvest vegetative growth
PHSD	Post harvest semi-dormant
PHD	Post harvest dormant

Table 3.8: Weeds present at pyrethrum data collection sites.

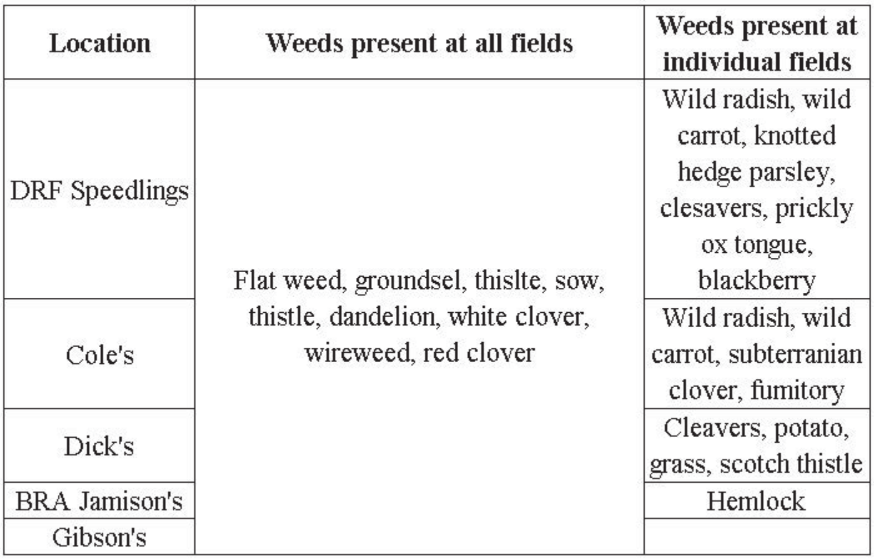

Location	Weeds present at all fields	Weeds present at individual fields
DRF Speedlings	Flat weed, groundsel, thislte, sow, thistle, dandelion, white clover, wireweed, red clover	Wild radish, wild carrot, knotted hedge parsley, clesavers, prickly ox tongue, blackberry
Cole's		Wild radish, wild carrot, subterranian clover, fumitory
Dick's		Cleavers, potato, grass, scotch thistle
BRA Jamison's		Hemlock
Gibson's		

Species identification of individual weeds in collected video was not performed as the image analysis was only required to discriminate two categories being pyrethrum and not-pyrethrum.

Chapter 5

Segmentation and the DCSA

5.1 Introduction

Image segmentation is the separation of pixels in an image into segments or groups of pixels that are similar for the purpose of reducing the complexity of the image data for further analysis. In the case of weed identification, the groups of pixels are generally green (plant) pixels and background pixels. Typical uses of image segmentation are to locate objects in an image. For example if the requirement is to follow a red ball across a video sequence, the segmentation aim would be to separate all the red pixels from non red pixels so the red pixels can be further analysed. Segmentation can be simple (e.g. thresholding pixel colour levels against a known value, as would be useful in locating the red ball) or complex (using texture descriptors to locate homogeneous regions in an image).

This chapter sets out the process of evaluation of common segmentation methods for use in weed spot spraying and the factors affecting the performance of each. An original segmentation technique called the Depth Colour Segmentation Algorithm (DCSA), combining colour and depth data, has been developed in this research which improves upon the results of the common segmentation methods. The operation of the DCSA segmentation technique is detailed and evaluated for application to pyrethrum and sugarcane crops.

5.2 Evaluation of common segmentation techniques

The literature review (Chapter 2) included segmentation techniques and highlighted the factors affecting the performance of segmentation algorithms. In particular, significant negative factors are introduced by real-world conditions. Real-world conditions provide a myriad of variations in stubble cover, plants at differing growth stages, and different levels of plant health. The literature review concluded that occlusion and illumination are significant obstacles encountered in the real-world environment and are difficult to contend with.

5.2.1 Occlusion and illumination effects on segmentation

Slaughter et al. (2008) stated that occlusion is the most significant factor to overcome for a real-time, real-world weed identification system. The literature review, surveyed techniques to separate occluding leaves by segmentation and feature extraction/classification. Using feature extraction and classification to overcome occlusion can involve an extra step in the image analysis process as further feature extraction and classification may be required after segmentation for plant identification. The extra feature extraction and classification step can add computational time and make real-time analysis difficult. Therefore, occlusion is preferably addressed as part of the segmentation analysis, if possible.

Segmentation techniques reviewed in the literature (and evaluated below, Section 5.2.3) were not satisfactorily able to segment plants with occlusions. The introduction of depth data, with colour, improves segmentation of occluded leaves as shown by Seatovic et al. (2008), Wallenberg et al. (2011) and Chene et al. (2012). However, the algorithms did not have real-time capabilities and the results were not developed for real-world conditions. Real-time capable depth segmentation techniques, evaluated below (Section 5.2.4), gave unsatisfactory results.

The literature review showed that illumination variation introduced a significant amount of error into segmentation. One means of overcoming the variation in illumination was the use of a light restricting cover over the viewing area. Therefore, research in this thesis used a light restricting cover over the viewing area to overcome the daylight illumination (Section 3.2).

5.2.2 Evaluation methodology

Zhang et al. (2008) state that evaluation of image segmentation techniques can be achieved by the following methods:

1. **Subjectively**. Subjective evaluation is where a person visually compares the segmented images to non-segmented originals and determines a segmentation quality.
2. **Supervised objective evaluation.** Supervised objective evaluation is tied to specific applications and is where the segmented image is compared to a manually ground-truthed reference image. This technique can be automated.
3. **Unsupervised objective evaluation.** Unsupervised objective evaluation is where the quality of segmentation is determined from the segmented image only, i.e. there is no ground-truthed reference image.

Subjective analysis is most commonly used with supervised objective evaluation also being common but unsupervised objective evaluation rarely used (Zhang et al. 2008). This present research uses subjective and supervised objective evaluation methods to evaluate different segmentation algorithms.

5.2.3 Evaluation of common colour segmentation techniques on real-world sample images

The initial step in the development of a weed detection algorithm in this research was to determine if common segmentation techniques performed satisfactorily enough to be used as a foundation for development of the image analysis system. To this end, a sample of images was collected which contained pyrethrum and grass-like plants. The sample images were then analysed by common segmentation techniques. The evaluation of the common segmentation techniques has been split into two groups based upon the computational complexity of the algorithm. These are ‘computationally expensive’ (requiring significant CPU resources); and ‘computationally inexpensive’ (requiring few CPU resources).

5.2.3.1 Computationally expensive colour-based segmentation techniques

A freeware segmentation tool called BVwin, distributed by Trolltech AS Norway[1], was used to undertake the evaluation of common segmentation techniques. BVwin provided a visual output and could segment in the RGB, HSV and greyscale colour spaces whilst processing the images with the segmentation techniques of region growing (Fan et al. 2001) (Figure 5.1(b)), colour structure code (Hartmann 1987) (Figure 5.2(a)) and split and merge (Haralick & Shapiro 1985) (Figure 5.2(b)). The visual result comprised identifying pixels of the same component with the same colour in the resultant image. The execution time for each segmentation technique was estimated at 1 s, 1.5 s, and 3 s respectively. Timing was taken with

---

[1] `http://www.codeforge.com/read/242483/license.txt__html`

a stopwatch whilst running the application on a dual core, 2.7 GHz computer, averaged over 10 executions of each technique.

The assessment of the segmentation techniques evaluated three different regions within the image that contained different plant features for segmentation as set out in column 1, Table 5.1. Columns 2 to 5 detail the segmentation technique and segmentation performance relative to the regions in column 1.

The results in Table 5.1 highlight the errors introduced to common segmentation techniques from illumination and occlusion, with none of the assessed algorithms performing well in all three labeled regions. Region growing performed the worst not segmenting any of the three Regions Of Interest (ROI) correctly. Split and merge performed the best with one correctly segmented ROI and one partially segmented ROI. All three techniques failed to segment the occluded plants (Label 1 in Figure 5.1).

Table 5.1: Results of common segmentation technique applied to three labeled regions of Figure 5.1(a).

Plant description	**Region in Figure 5.1(a)**	**Region growing**	**Colour structure code**	**Split and merge**
sorghum occluding pyrethrum	1	occluded	merged with ground	mis-labelled
isolated pyrethrum	2	merged with ground	correctly segmented	correctly segmented
partial pyrethrum plant	3	merged with ground	merged with ground	partially segmented

(a) Original colour image of plants. Table 5.1 outline the red ellipses.

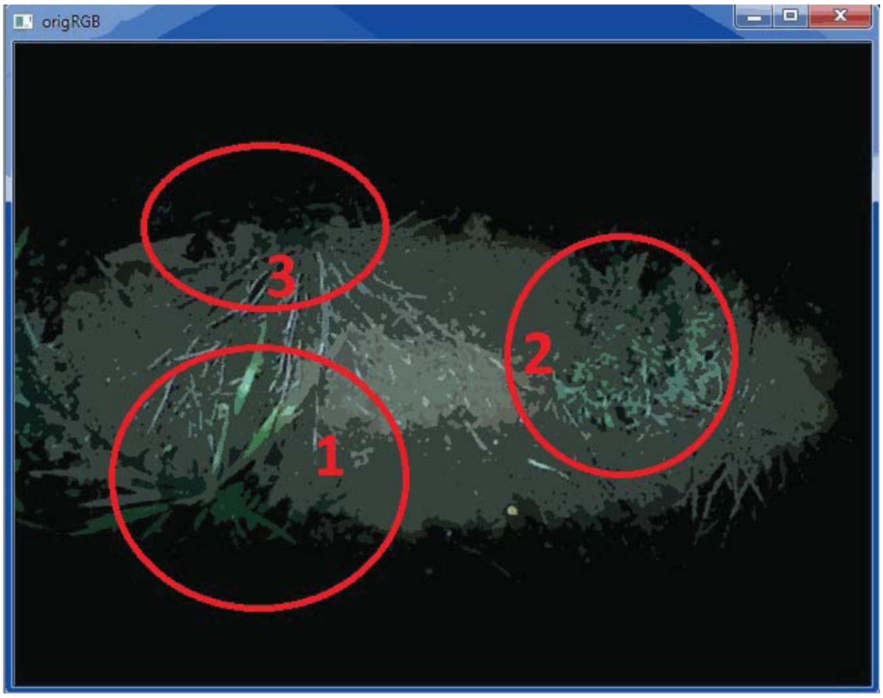

(b) Region growing segmentation technique applied to a colour image 5.1(a).

Figure 5.1: Image sequence showing the original image and results of BVWin segmentation implementations.

(a) Colour structure code segmentation technique applied to a colour image 5.1(a).

(b) Split and merge segmentation technique applied to a colour image 5.1(a).

Figure 5.2: Image sequence showing the results of BVWin segmentation implementations.

5.2.3.2 Computationally inexpensive colour-based segmentation techniques

Computationally inexpensive segmentation algorithms are simple in operation, and beneficial in a real-time systems but can perform poorly in segmenting occluded leaves. Binarisation is an example of a computationally inexpensive segmentation technique. Figure 5.3 is a Binarised Segmentation Technique (BST) ($G > R$ and $G > B$) applied to Figure 5.1(a) used by Sabeenian & Palanisamy (2009) and McCarthy et al. (2012). The pixels inside the yellow circle in Figure 5.3 show that 2D colour based segmentation techniques are impacted greatly by occlusion as the white pixels of the two different plants appear connected.

In order to address the occlusion of the plants in the yellow circle the BST technique was modified so that a multiplier was applied to the red and blue channels to reduce sensitivity to green by 10%. The results, Figure 5.4 show the occlusion was reduced by decreasing sensitivity to green, but green plant material was also lost in the process. Features present in Figures 5.3 and 5.4 are metamerism[2] and noise which can both create false positives in the segmented images, specifically in the darker edges of the lit area and this can be a significant source of error for a colour only system.

Figure 5.4 displays sensitivity to illumination and occlusion which was also found in the computationally expensive segmentation techniques. Sensitivity to illumination is displayed by the pattern of false triggers in the BST images (Figures 5.3 and 5.4). Figure 5.1(a) is the colour image associated with the BST images and the centre of the images as well lit. Small false triggers appear in the BST images where the brightness of the light is reduced in the colour image and soil is mistaken as green.

---

[2] Metamerism is the incorrect representation of colour by a set of RGB pixels of an object which has different spectral power distributions (Fairchild et al. 2014).

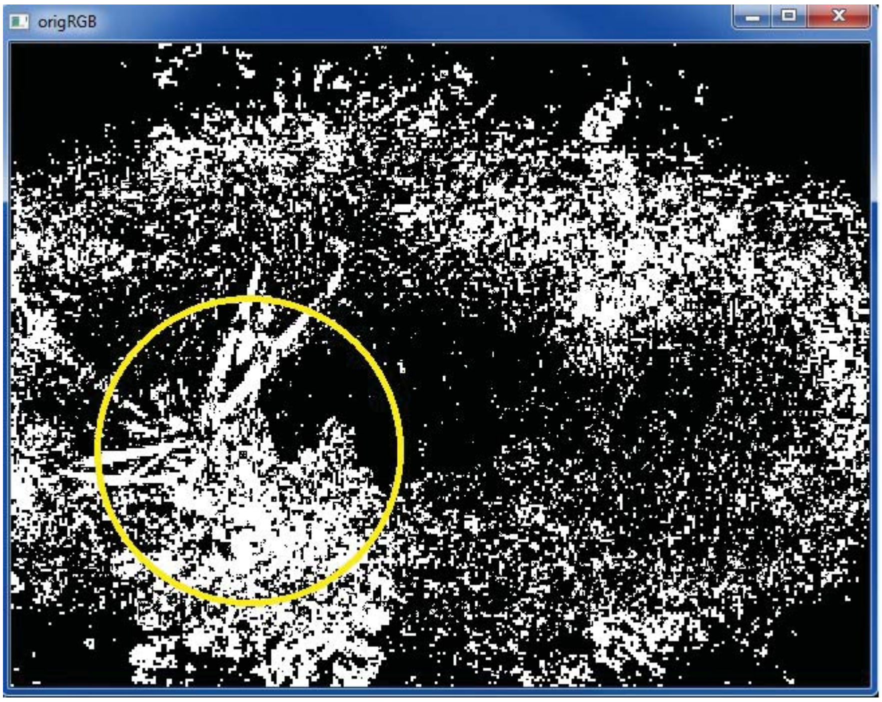

Figure 5.3: Binarised image of Figure 5.1(a) using a BST. Green leaves and dark areas of image are segmented as vegetation.

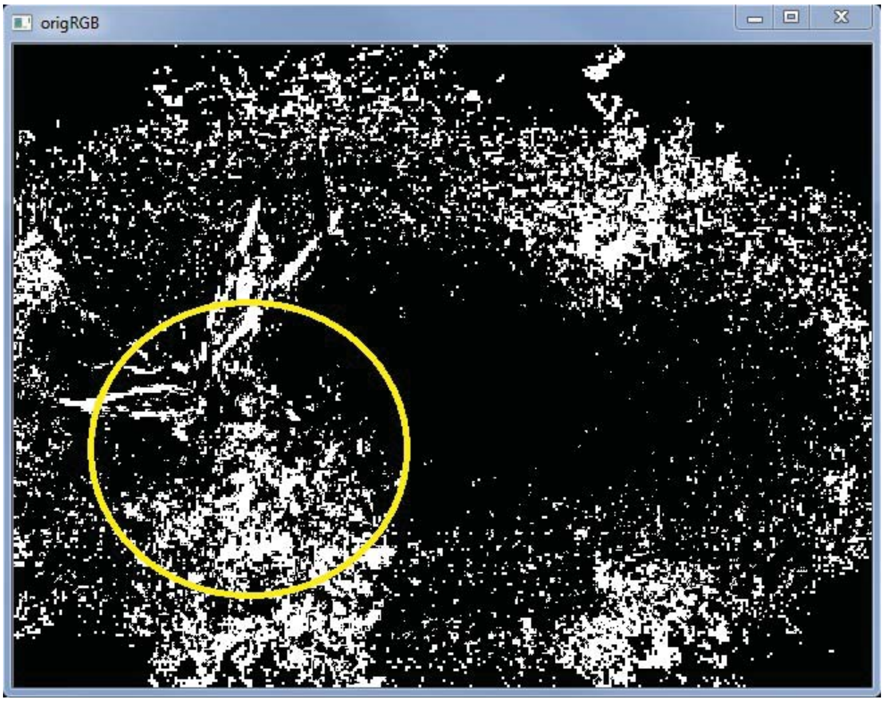

Figure 5.4: Binarised image of Figure 5.1(a) using a BST where R and B are reduced by 10% to lessen the false positives in the image.

5.2.4 Evaluation of depth segmentation techniques on real-world sample images

Section 5.2.1 stated that the depth segmentation techniques and depth/colour segmentation techniques in the literature review (Chapter 2) were not effective in real-time systems due to the computation time required and adaptability to real-world situations. Two depth segmentation techniques have been identified as viable for real-time operation, namely connected components[3] functions and thresholding.

5.2.4.1 Thresholding

Thresholding separates the pixels of an image into groups of pixels that are above the threshold depth value and those pixel values below the threshold depth value. Accurate thresholding of plant material from the ground would be difficult to achieve in a real-world setting at commercial groundspeeds as the threshold value would have to accommodate the continual deviation in height of the camera. The spray boom section holding the camera will deviate in height (hence so will the camera) caused by the roughness of the ground, groundspeed, tyre pressures and ground undulation. Therefore grasses that run along the ground and plants that are prone to the ground would be intermittently grouped with those pixels below the threshold value and disregarded. Alternatively ground and stubble would be intermittently grouped with those pixels above the threshold and kept for analysis when they should not be.

---

[3] A connected components analysis (also known as floodfill or seedfill) is used to group pixels of similar intensities or within a set variation of intensities, into a contiguous shape (component) that all have the same unique label (Bradski & Kaehler 2008).

5.2.4.2 Connected component functions

Performing a connected components function on the depth image groups those pixels with values similar (or within an allowable deviation) to their neighbouring pixels. The connected componenst result is shown in Figure 5.6 with each component having a unique colour label. Figure 5.5 is the colour image associated with the resultant connected components depth image. Comparing Figures 5.5 and 5.6 highlights the difficulty the connected component segmentation technique had in segmenting the plant from ground. The segmentation error is seen in Figure 5.6 with the pyrethrum plant (red ellipse labeled '1'), merging with the ground (red ellipse labeled '2') where the stubble and plant are combined.

Figure 5.5: Colour image of a pyrethrum plant and a sorghum plant.

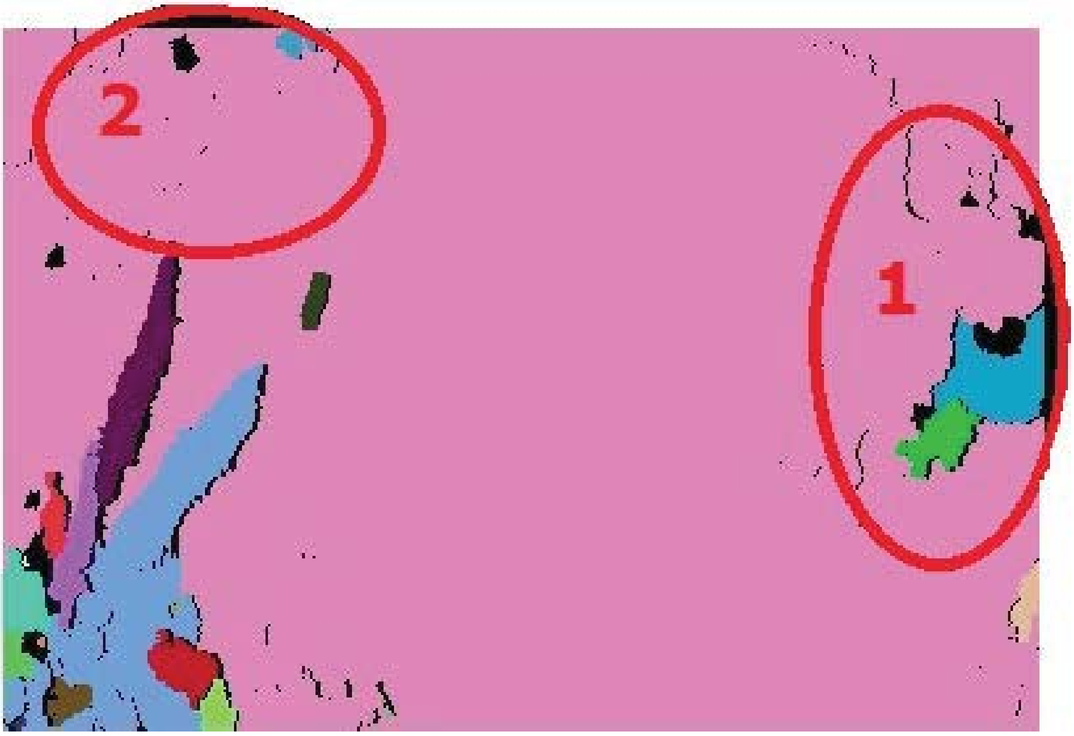

Figure 5.6: A connected components applied to the depth image associated with Figure 5.5 where each component is assigned a unique colour. Errors are apparent in red ellipse one where ground and plant merge and in red ellipse two where stubble and plant merge.

5.2.5 Summary of common segmentation techniques

Visual inspection of the results in Figure 5.2 to Figure 5.6 reveal that existing segmentation techniques do not robustly segment different plants in real-world settings due to occlusion, illumination and low lying plant positions. These results are supported by the literature review. Existing segmentation techniques group occluding plant material together in the one object which creates errors in feature extraction and classification. Therefore there is a need for new segmentation techniques to be developed that can operate in a real-world environment and in real-time.

5.3 Development of the Depth and Colour Segmentation Algorithm (DCSA)

The aim of the Depth and Colour Segmentation Algorithm (DCSA) is to be a real-time segmentation function that can segment individual plant components (in fallow or crop) from other plants and foreign objects (e.g. stubble and rocks) with a high level of accuracy when occluded in minimum and no-till situations. In the following sections, the DCSA is shown to be successful segmenting weeds in fallow situations and in crops that have differing growth stages and therefore varying leaf shape, height and colour.

5.3.1 The DCSA as a modified connected component algorithm

The DCSA segments an image into separate components (leaves) within the image, based on their colour and depth connectedness. The DCSA achieves segmentation using a modified connected components analysis. The operation of a standard connected component algorithm is as follows.

Connected component algorithms start by locating a ‘seed’ position (starting point) in the image and then evaluating the pixels around it. The pixels around the seed pixel that are within given tolerances of similarity are labeled the same as the seed pixel. The analysis then moves its seed position onto one of the newly labeled pixels and repeats this process; the repeating continues until there are no new pixels that are within given tolerances remaining in the image. The group of labeled pixels is now a completed individual object or ‘component’.

The connectedness of the pixels in a connected components analysis can be either four way connectivity or eight way connectivity (Bradski & Kaehler 2008) as set out in Figure 5.7. In the four way connectivity method, the connected components analysis evaluates the pixels numbered 1 to 4 individually against the seed pixel's value. In the eight way connectivity method, the connected components analysis evaluates the pixels numbered 1 to 8 individually against the seed pixel's value.

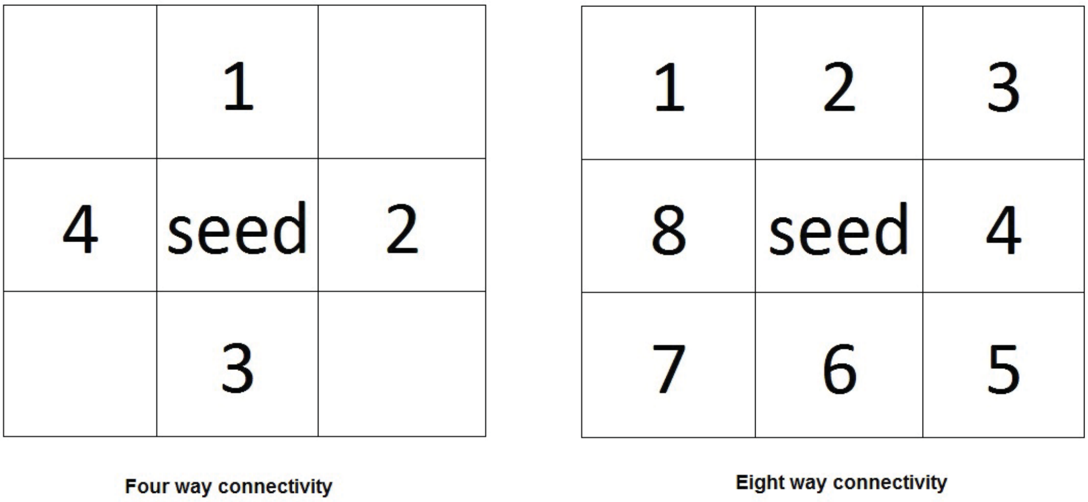

Four way connectivity

Eight way connectivity

Figure 5.7: Four and eight way connectivity diagrams indicating the directions from the seed pixel that are evaluated.

5.3.2 Colour and depth connectivity with flag-driven configurability

The DCSA modifies a standard connected components algorithm to incorporate two data streams, being colour and depth. The software implementation of the modified connected component is 'flag-driven' to allow for flexibility and selectability. Flag-driven means that the DCSA will apply certain additional segmentation analysis to the image dependent upon the analysis capability being enabled or disabled by a flag. A flag is either set (i.e. 1 or TRUE) or cleared (i.e. 0 or FALSE). The DCSA configuration flags can be an argument parsed to the DCSA function when called in the program or a shared variable between the DCSA function and the overall image analysis process. The configurability of the DCSA provides flexibility to the user as the flags can be modified either at the start of the overall real-time spot spray operation (i.e. in a new paddock) or 'on-the-go' (between frames).

The DCSA's operation is outlined below in Section 5.3.3 by means of an example. The concept requires depth and image data (RGB, greyscale or spectral-other than RGB) as inputs and the resultant segmented components are saved for further analysis.

5.3.3 An example of DCSA operation

5.3.3.1 DCSA scan

The images used in the DCSA example are from a Kinect® camera system described below with the co-ordinate system defined in Figure 5.8:

- A greyscale image (eight data bits per pixel) representing depth by the pixel intensity. The higher the intensity of the pixel the closer the object in the image is to the camera.
- A colour image (three channels (R,G,B) each channel with eight data bits per pixel) of the same scene as the greyscale image with the pixels of the colour and greyscale images mapped, i.e the pixels representing the scene on one image directly relate to the pixel positions of the same scene on the other image.

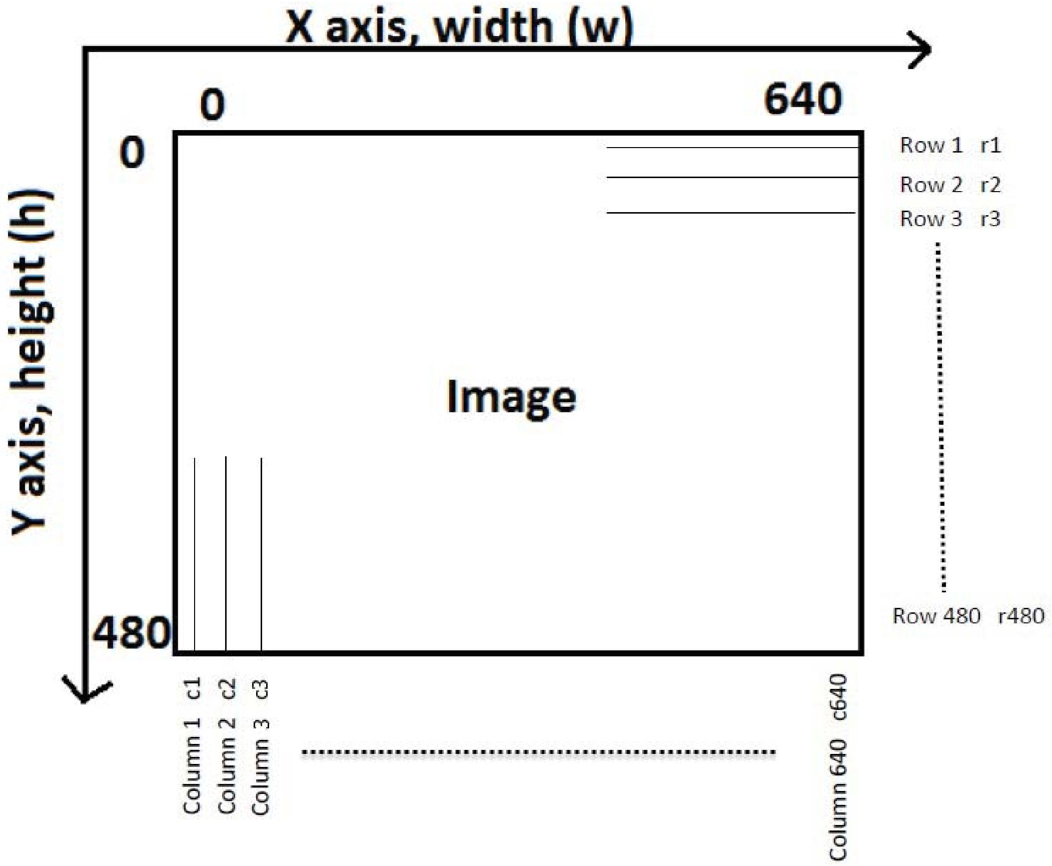

Figure 5.8: Co-ordinate system for the DCSA, with image size 640 × 480 pixels.

The DCSA initiates by scanning the depth image to locate a seed pixel (i.e. somewhere to start the component). The seed pixel scan starts at column *c1*, row *r1* and increments across the columns of *r1*, from *c1* to *c640*. The row then advances to *r2* and the scan from *c1* to *c640* repeats (Figure 5.8). This continues until *r480* is reached. The seed pixel scan direction can be reversed, i.e. incrementing the rows and advance along the columns without affecting the operation of the process.

5.3.3.2 DCSA connectivity analysis

The seed pixel scan is searching for a pixel that has a value (height) above a user set value (height *H2* in Figure 5.9). When a pixel above *H2* is found, the pixel is labeled with a unique identifier and follows 4 way connectivity to find pixels of similar value or within a range of values and label them with the same identifier. As outlined in Section 5.3.1, the labeled pixels are then used as seed pixels to look for neighbouring pixels that have an value within a threshold amount, $connected_{\mathrm{thresh}}$, of the value of the seed pixel.

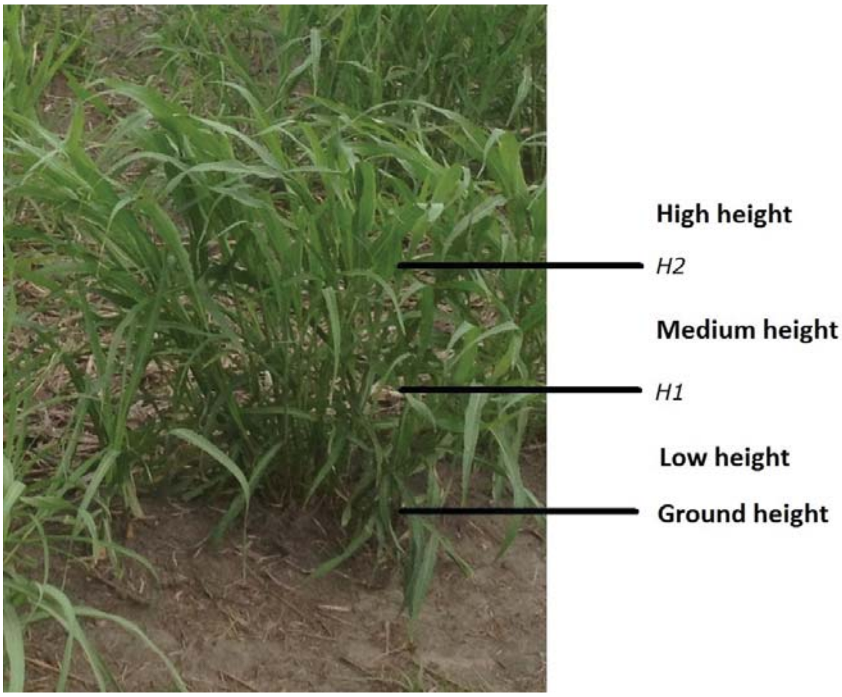

Figure 5.9: Plant height definition for the DCSA classification criteria relative to a sugarcane plant. *H1* and *H2* are user-chosen threshold points for determining low, medium and high areas of a plant.

For example, if $connected_{\mathrm{thresh}}$ is set to 3, then a neighbouring pixel with an value within ± 3 would be accepted as connected to the seed pixel and relabeled to the unique identifier of the seed pixel. The connected pixel is subsequently used as a seed pixel in the continuing connectedness search. If the value of the neighbouring pixel is greater than ± 3 from the seed pixel, the neighbouring pixel is considered to be an edge and not part of the same component.

5.3.3.3 DCSA continuing process

The connectedness process is continued until no new pixels are found to add to the component. Once the component is complete (i.e. no new connected pixels have been found), the seed pixel scan of the image continues searching for a pixel with intensity greater than *H2* and not already labeled. This process continues until there are no pixels left unlabeled greater than *H2* in the depth image. The seed pixel scan position then resets to *r1*, *c1* and a second scan is initiated.

The second seed pixel scan searches for a pixel that has not previously been labeled, or that is not attributed to noise (indicated by a pixel value of 255, determined in the pre-processing Section 4.3). When a pixel is found, all connected pixels are identified and the image seed pixel scan continues and repeats the process of forming new components for each unallocated seed pixel. At the end of this process all plant pixels are labeled.

5.3.3.4 DCSA additional analyses

As previously noted, a novel additional capability of the DCSA is having connectivity in a component based on connectivity in the depth image and in the colour image. Depth and colour connectivity require the connected component algorithm to search the same pixel position in the colour image as well as searching the depth image pixel position (enabled by setting the flag GREENCHECK). A pixel is only added to the component in the depth image if the corresponding pixel in the colour image is green (i.e $G > R$ and $G > B$), otherwise the pixel is left unlabeled as a non-plant pixel.

Additional flags in the DCSA provide the following analysis capabilities:

1. **Tolerance of image alignment errors in pre-processing (Section 4.3.2.3).**

   Tolerance of image alignment is achieved by allowing the high portions of the leaf/component in the image to be connected in the depth image only and partially connected in the colour image. To enable tolerance in the DCSA, the flag ALIGNMENT_ERROR must be set and the green pixels in the component above *H1* are counted (to determine total green pixels in the component) but the pixel's connectedness in the component above *H1* is only determined by the depth pixel's connectedness. Below *H1* operates as per GREENCHECK flag description previously described in this section. When the component is complete and being sorted into the retained or deleted images (Section 5.3.3.5 below), the percentage green of a component above *H1* determined by total green pixels divided total component size in pixels, is thresholded with $percentage\_green_{\text{thresh}}$. The value of $percentage\_green_{\text{thresh}}$ provides the tolerance in the alignment of the depth and colour images, i.e if $percentage\_green_{\text{thresh}}$ is set to 90% then the component has little tolerance to misalignment and if set to 50% a significant amount of misalignment tolerance is allowed.

2. **The flexibility to limit the total variation of intensities in the component.**

   A small limit of intensity variation in the component (e.g. 1-5 in pixel value) creates segmented components of depth slices where the same tall leaf is separated into multiple neighbouring components with graduating heights. Whereas a large limit of intensity variation in the component (e.g. 50 in pixel value) provides a segmentation where the component has potential to start at a large height (e.g. the top leaves of a sugarcane plant) and generate a single component that traverses a tall leaf down to a low height or even the ground level. Limiting the variation in intensity is achieved by setting a flag MAXLENGTH and then a value $variation_{\text{max}}$ imposes a limit on the absolute variation between a neighbouring pixel's intensity and the first (original) seed pixel value of ± $variation_{\max}$.

3. **Statistics can be determined and recorded for each component.** The colour statistics maximum hue, minimum hue, average hue and hue variance in the component are recorded when the COLOURSTATS flag is set. Maximum height, minimum height, average height and height variance statistics are recorded when the DEPTHSTATS flag is set. Bounding box position and total size (in pixels) of the component are always recorded. These statistics can be used as features in the feature extraction process.

5.3.3.5 DCSA component sorting

The extracted components are sorted into scratch images (temporary images created for internal analysis) defined as ‘retained image’, ‘deleted image’ or ‘debug image’ as follows.

- The retained image contains segmented components used further in the analysis process.
- The deleted image contains segmented components that are not required for further analysis.
- The debug image is a visual representation of all components for debug purposes during algorithm development.

The sorting of the components into the retained and deleted images is determined by flag settings in the DCSA.

Table 5.2 contains pseudo code to sort the DCSA components into retained and deleted images.

Table 5.2: Pseudo code to sort components as retained or deleted.

Step	Operation
1	**if** the component size (in total pixels) is less than a threshold size ($total\_size_{\text{thresh}}$), delete the component. **else** continue
2	**if** the REDUCE_LINES flag is set (TRUE) then run the reduce lines function on the component. Reduce lines divides the bounding box width by height. **if** the result is $> RLMin_{\text{thresh}}$ and $< RLMax_{\text{thresh}}$ (0.5 and 2 respectively) the component is deleted, otherwise the total component size is divided by the bounding box area and **if** the result is $< RLSize_{\text{thresh}}$ (0.3) the component is deleted.
3	**if** the average depth of the component is less than or equal to *H1* and the KEEP_LOW flag is set (TRUE), retain the component.
4	**else if** the average depth is greater than *H1* and less than *H2* and the KEEP_MEDIUM flag is set (TRUE), retain the component.
5	**else if** the average depth is greater than *H2* and the KEEP_HIGH flag is set (TRUE) and ALIGNMENT_ERROR flag is set (false) then retain the component.
6	**else if** the average depth is greater than *H2* and the KEEP_HIGH flag is set (TRUE) and if the percentage of green pixels in the component is greater than a percentage ($percentage\_green_{\text{thresh}}$) and the ALIGNMENT_ERROR flag is set (TRUE) then retain the component. The value of $percentage\_green_{\text{thresh}}$ provides leniency for pre-processing alignment errors in Section 4.3.2.3.
7	**else** delete the component.

5.4 DCSA features and limitations

5.4.1 DCSA features

5.4.1.1 Segmentation accuracy

If depth data alone is considered when identifying components, a component that is connected to the ground (e.g. grass leaves, plant stems and clovers) can be labeled with the ground as a contiguous object (e.g. Figure 5.6). Addition of a criterion that requires a pixel to be green for that pixel to be added to the component prevents the ground from being grouped with a leaf and will correctly segment the component. Likewise depth data alone may also erroneously segment stubble and foreign objects with plant pixels. Again the addition of the criterion for a pixel to be green avoids this error.

Typical colour-only segmentation techniques (e.g. ‘excess green’), have limited capacity to segment occluded plants and leaves. The addition of a depth criterion, that requires the component to be contiguous in the depth image as well as the colour image, allows individual plants to be segmented as long as the edge of the component can be found in either the colour or depth data.

5.4.1.2 Sorting capability

The sorting capability of the DCSA is able to reduce the data requiring further analysis in the feature extraction and classification area. The DCSA achieves this by collecting information about the component as it is being formed and this information is then used to sort components into components requiring feature extraction and classification and components of no further interest.

5.4.2 Known DCSA limitations

5.4.2.1 Component merging

As previously mentioned the DCSA can find the edge of a component when there is an edge in either the depth or the colour image. Accordingly if there are no edges between the components in the depth and colour image, the two components will be merged as one. Figure 5.10 shows a sugarcane plant where two leaves of sugarcane exhibit the same colour and cross each other at the same height, therefore merging into one component.

Figure 5.10: Image showing the depth image, segmented image and colour image associated with each other. The white ellipse highlights two separate leaf components that overlapped in colour and height and therefore are the one component.

5.4.2.2 Component splitting

The DCSA segments the images into individual leaves. Plants that grow in patches with indistinct height attributes (e.g. couch grass and vines) can be segmented into a number of components instead of a single component as shown in Figure 5.11.

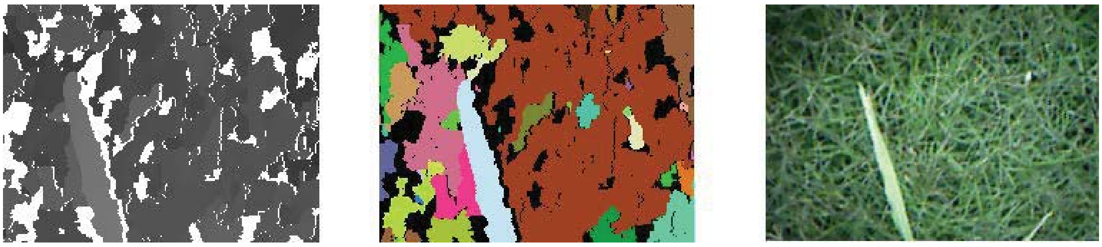

Figure 5.11: (Left to right) depth image, DCSA segmented image and original colour image associated with each other for couch grass and single sugarcane leaf. As shown (centre) numerous components represent the couch grass.

5.5 Field trials for DCSA evaluation

5.5.1 Comparison of results for occlusion

An objective of this research is to show that combining colour and depth provides a more robust result than either depth or colour on their own (Section 1.3). Evaluation of common colour segmentation techniques (Section 5.2.3) highlighted that complex segmentation techniques of region growing, colour code, and split and merge could segment plants in a small degree, but were not real-time operable. Simple segmentation techniques of excess green and binarisation were real-time ready. Evaluation of depth segmentation techniques (Section 5.2.4) drew attention to the inability of complex depth image segmentation techniques to operate in real-time and the real-world inadequacy of two techniques (thresholding and connected components) that were real-time operable.

Therefore, in the following field evaluations the DCSA was compared to a Binarised Segmentation Technique (BST) for colour image segmentation and no depth segmentation technique has been compared.

5.5.2 Collection of evaluation data

Additions were made to the DCSA analysis software to provide data collection and other analysis for DCSA evaluation. The modifications were:

1. to apply a BST to the colour image separately from the DCSA;
2. to record the number of segmented components and the number of pixels in the components for the DCSA and BST before sorting the components into the deleted and retained images;
3. to record the number of segmented components and the number of pixels in the components for the DCSA and BST after sorting the components into the deleted and retained images; and
4. to manually assess and record the status of occlusion in the colour image and status of occlusion before and after the segmentation techniques were applied.

Field trials were then conducted in sugarcane and pyrethrum as set out below.

5.6 DCSA evaluation in sugarcane

5.6.1 Sugarcane and guinea grass growth attributes

Sugarcane is a perennial crop which can have up to 5 ratoons (seasons of re-growth) in the Australian sugarcane cropping areas and often has a ground cover of the previous season's harvest residue (trash) which is referred to as a trash blanket (Appendix C.1). Guinea grass was the target weed for discrimination from sugarcane identified by the Sugar Research Development Corporation (SRDC).

reducing the amount of data requiring further analysis shortens processing time further in the image analysis process. The reduction in processing time is due to the reduction in pixel data requiring further analysis in feature extraction and classification (Chapter 6).

5.6.3.5 Summary of results for the DCSA application in sugarcane

The results for the evaluation of the DCSA in sugarcane shows that for the test dataset:

1. the DCSA sorting can reduce the pixels requiring feature extraction and classification by up to 84% (76% compared to BST) and components by up to 54%;
2. the DCSA is robust with a high accuracy when occluded ($> 99\%$); and
3. the DCSA can operate in real-time at less than 23 ms per frame.

5.7 DCSA evaluation in pyrethrum

5.7.1 Pyrethrum growth attributes

In the DCSA's application to pyrethrum, the requirement was to identify the pyrethrum crop and to spray all plant material not pyrethrum. Weed control for pyrethrum is undertaken from harvest in January through to May. Pyrethrum grows to a height of approximately 0.3 m and a minimum diameter of approximately 0.2 m by the end of summer, and then lays dormant over winter. In general the plant density is adequate to create a continuous row of pyrethrum, with the occasional small break in the row caused by harvest damage.

5.7.2 DCSA setup and operation in pyrethrum

5.7.2.1 DCSA setup parameters for pyrethrum operation

The parameters for sorting component of the DCSA was set to retain those components which:

- have a maximum height between *H1* and *H2*;
- are green; and
- are above a threshold size.

The parameters required to achieve the retained image results are set out in Table 5.8.

5.7.2.2 DCSA operation in pyrethrum

The DCSA operated on pyrethrum images by deleting all components whose average depth was less than *H1* and greater than *H2*. The DCSA deleted all components that were below the size of minimum pyrethrum ($total\_size_{\text{thresh}}$). Therefore if a plant had a height between *H1* and *H2* and was larger than what was considered the minimum size for pyrethrum at that particular growth stage, it was retained.

Table 5.8: DCSA setting identified in Section 5.3.3 for pyrethrum. X signifies 'don't care'.

Setting / flag	**Value**
DEPTHSTATS	true
COLOURSTATS	false
$total\_size_{\text{thresh}}$	1000
GREENCHECK	true
ALIGNMENT_ERROR	X
$percentage\_green_{\text{thresh}}$	X
REDUCE_LINES	false
$RLMin_{\text{thresh}}$	X
$RLMax_{\text{thresh}}$	X
$RLSize_{\text{thresh}}$	X
KEEP_LOW	false
KEEP_MEDIUM	true
KEEP_HIGH	false
H1	20
H2	40
MAXLENGTH	true
$variation_{\text{max}}$	20
$connected_{\text{thresh}}$	3

5.7.3 Results and discussion for the DCSA technique used in pyrethrum

5.7.3.1 Evaluation data and setup

To evaluate the segmentation techniques, 500 sequential frames of real-world data taken from the DRF-Speedlings site on the 11-04-13 (Table 3.3) were analysed. This dataset was considered typical of pyrethrum crop conditions and contained examples of weeds found only at the DRF-Speedlings site and also weeds that were common to all fields (Table 3.8). The growth stage at April was in the weed control timing range (January to May). From visual inspection, the pyrethrum varied in height between 0.25 m and 0.35 m therefore *H1* was set to 0.2 m, *H2* was set to 0.45 m and the minimum size threshold for component size was set to 0.15 m. All settings and flags are shown in Table 5.8.

5.7.3.2 Results for occlusion in pyrethrum

There were 25 occluded pyrethrum plants in Table 5.9 and 53 non-occluded pyrethrum plants. Table 5.9 highlights the DCSA segmented 100% of the occluded pyretrhum plants from weeds whereas the BST was not able to segment any of the occluded pyrethrum plants from the weeds.

Table 5.9: Statistics for the occlusion tolerance of the DCSA compared to a 2D (colour) binarised segmentation technique in pyrethrum.

Total number of weeds	78
Number of weeds occluded in video data	25
Number and percentage of weeds occluded after BST	25 (100%)
Number and percentage of weeds occluded after DCSA	0 (0%)

5.7.3.3 Results for sorting in pyrethrum

The result for the sorting capability of the DCSA is set out in Tables 5.10 and 5.11, which show that the DCSA reduced the number of components for feature extraction and classification by 49% and the pixels by 55%. The standard deviation of the components and pixels in Tables 5.10 and 5.11 indicate a high amount of variation from frame to frame which can be attributed to the changing amount of plant material between the frames.

Table 5.10: Reduction of components in the DCSA in pyrethrum; statistics of experimental results.

Statistic	**Average**	**std deviation**
Number of depth components per image	85	37
Number of depth components after segmentation	44	25
Percentage reduction in the number of depth components	49%	-
Computation time of DCSA in ms	6.6	1

Table 5.11: Reduction of pixels in the DCSA in pyrethrum.

Statistic	**Average**	**std deviation**
Depth component pixels per frame before segmentation	65191	16266
Depth component pixels per frame after segmentation	29575	16746
Percentage reduction in depth component pixels	55%	-

Table 5.12 displays a reduction in pixels after DCSA sorting compared to the BST of 13%.

Table 5.12: Reduction of pixels from the binarised segmentation technique in pyrethrum.

Statistic	**Average**	**std deviation**
Number of pixels per frame after BST segmentation	33875	13578
Reduction retained depth pixels compared BST pixels	13%	-
Computation time of BST in ms	1.2	0.35

5.7.3.4 Real-time application to pyrethrum

Table 5.9 show that the DCSA execution time was less than 10 ms (average time plus 3 standard deviations, i.e. 99% of data and would exclude outliers) which is well within the realms of real-time system requirements. The BST took less than 2.5 ms which is also well within the realms of real-time system requirements.

5.7.3.5 Summary of results for the DCSA application in pyrethrum

The results for the evaluation of the DCSA in sugarcane shows that for the test dataset :

1. the DCSA can reduce the pixels requiring feature extraction and classification by up to 55% (compared to 13% for BST) and components by up to 49%;
2. the DCSA is robust with a high accuracy when occluded (100% observed); and
3. the DCSA can operate in real-time at less than 10 ms (average 6.6 ms) per frame.

5.8 Summary of Chapter 5 and results

This Chapter has:

- discussed the problems of occlusion and illumination for segmentation;
- evaluated common colour segmentation techniques as a means of segmenting plants—with unsatisfactory results;
- evaluated depth segmentation techniques—also with unsatisfactory results;
- highlighted the need for a new segmentation technique;
- described the development and operation of a novel, new segmentation technique (DCSA) which combines colour and depth in real-time. Components

formed by the DCSA have a connectedness defined by the degree of similarity of pixel colour and depth; and

- evaluated the DCSA for use in sugarcane and pyrethrum.

The results for the DCSA technique show that the addition of colour and the depth data aids significantly in identifying occlusion for segmentation of plants by being able to locate edges in either the depth or colour image. The DCSA can segment plant from stubble and, potentially, other foreign objects even when they are a similar height to the plant material. The evaluations in sugarcane and pyrethrum showed that the DCSA has a greater than 99% accuracy when occluded in the test data which satisfies the occlusion tolerance goal for the thesis. The BST was shown to have no occlusion tolerance capability.

The DCSA technique reduced the amount of data requiring further processing, compared to the BST, by 76% in sugarcane and 13% in pyrethrum. The variation in the results for this sorting capability between pyrethrum and sugarcane indicates that the DCSA technique offers greater benefits in crops were the crop is higher than the weed (or weed is higher than crop) and has different physical traits to the weed e.g. grass, broadleaf, clumping and leaf size. This was highlighted by an 84% reduction in pixels requiring feature extraction and classification in sugarcane and 55% in pyrethrum.

The execution time for the DCSA analysis of the pyrethrum was 10 ms, and 23 ms in the the sugarcane, which falls within the real-time requirements of this research. Therefore the DCSA:

- meets the aims of the thesis exhibiting a high accuracy when occluded;
- can segment individual leaves;
- improves execution time for feature extraction and classification by sorting the components; and
- fulfills real-time requirements.

discriminatory attribute because the pyrethrum grew to a height of 0.2 to 0.3 m in autumn (the primary weed control period) and then stayed at that height until spring. Spatially, the pyrethrum plants were grown in rows and the centre of the pyrethrum plant was centred on the centre of the row. Therefore, plants that were not centred on the row could be identified reliably as weed. To evaluate the usefulness of height and spatial position as a feature for pyrethrum identification, four algorithms were developed and are outlined below:

- **Spatial Position (SP).** The centre point of each pyrethrum plant should be the centre of the row so the spatial position of the plant component relative to the centre of the row of pyrethrum is used to aid pyrethrum identification.
- **Depth, Colour and Size (DCS).** This algorithm compares the depth, colour and size attributes of a plant component against a template.
- **Depth, Colour, Size and Spatial position (DCSS).** The DCS algorithm operation is aided by the SP concept of centralisation of the plant component over the row.
- **LBP and Depth (LBPD).** The LBPD adds a depth component to the window of LBP data and compares to a template.

6.6.1.1 Spatial position (SP) algorithm

The spatial position algorithm used the segmented binarised RGB image based on a Binarised Segmentation Technique (BST) ($G > R$ and $G > B$) and determined the presence of weed by the position of the segmented component in the image, relative to the centre of the image. It was assumed that the pyrethrum row would be centred in the image by a side-shift three-point linkage hitch (documented in Appendix G) guided by a vision guidance system which maintains the implement position centred over the row of plants.

The image was divided into five regions on the horizontal axis as shown in Figure 6.3 with region 3 being the central region. Five regions (odd number) were selected so that when a plant component was identified the plant component could be quickly identified as being ‘centred’ or ‘not centred’ depending upon the regions the Plant Component’s Bounding Box (PCBB) is contained in. If the PCBB appeared in more than the centre region and the total number of regions the PCBB was contained in was even, the PCBB was a weed, if odd the PCBB was pyrethrum. The centre region (3) was adjusted to the same width as the typical pyrethrum plant in the field so that if a pyrethrum plant was evaluated that was wider than the centre region, the PCBB would appear in the regions either side of the centre (odd number), centred on region three. Figure 6.3 displays the thirteen positions the weed PCBB may appear.

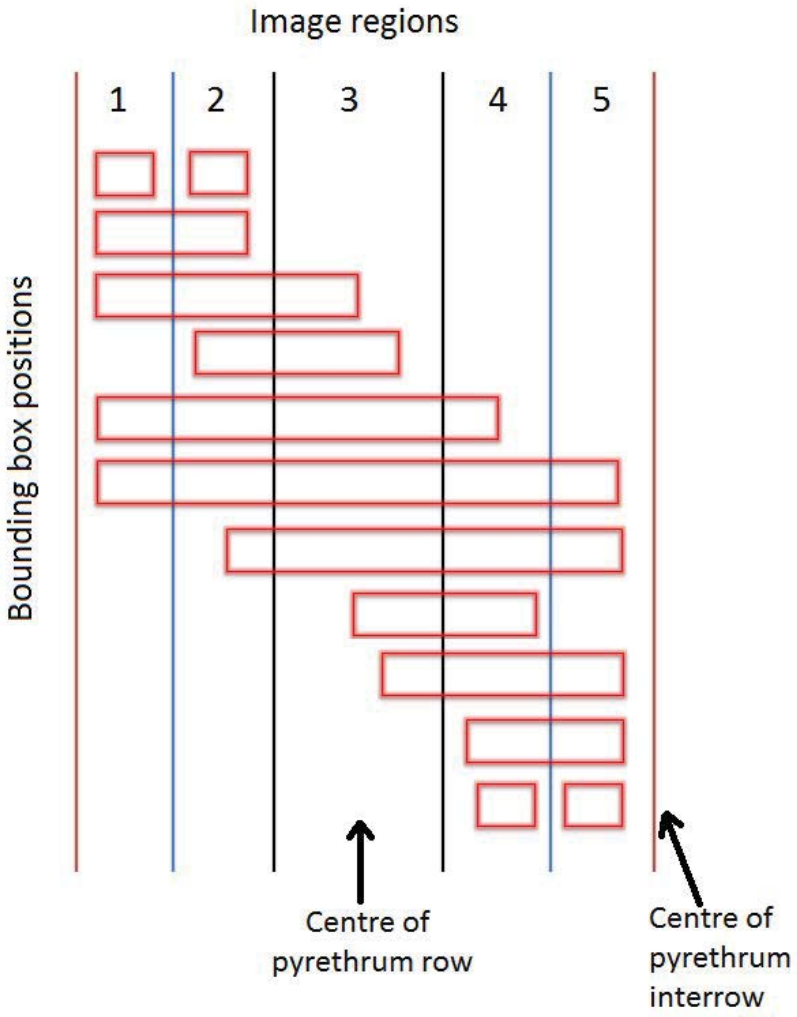

Figure 6.3: Bounding box positions used in spatial analysis that determine if a component is weed.

6.6.1.1.1 Evaluation of the SP algorithm

Evaluation of the spatial position algorithm was performed subjectively by colouring the image pixels to indicate where the system recognised a component as a weed, then comparing this to the same image 'untouched' i.e. initial colour image. The colours associated with the pixel components can be seen in Figures 6.4 and 6.5.

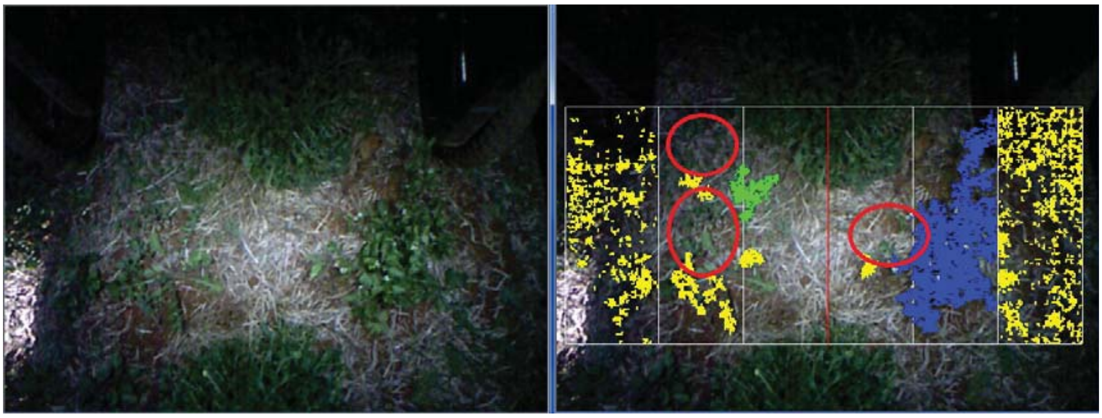

Figure 6.4: Results of the spatial segmentation method on small weeds highlighting misses from poor binarisation in the red ellipses. RGB image of pyrethrum and weeds on the left hand side. The images on the right hand side shows the weeds identified by shading the pixels were yellow, green, red, and blue depending on component size and position.

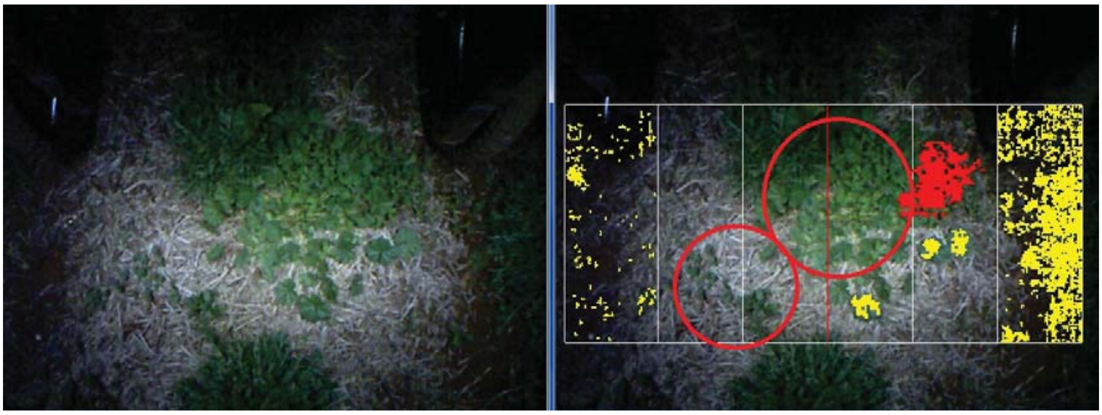

Figure 6.5: Results of the spatial segmentation method on large weeds highlighting misses from occlusion in the red ellipses. RGB image of pyrethrum and weeds on the left hand side. The images on the right hand side shows the weeds identified by shading the pixels yellow, red, blue and green depending on the weed's size and position. The red ellipses highlight areas where the weeds were not identified in Figure 6.3.

The colours were yellow, green, red, and blue depending on component size and position. Yellow was all weeds in regions one and five as well as the smallest weeds in regions two, three and four. Weeds coloured green, red and blue represented weeds from small to large respectively in regions 2, 3 and 4.

Poor algorithm performance was visually noticeable on images where there was occlusion, highlighted in the red ellipses in Figure 6.5. A further problem was misses of weeds that where determined to be errors in the BST of the colour image and are highlighted in the red ellipses in Figures 6.4 and 6.5. Modifying the BST created false triggers due to metamerism and poor image quality associated with cheaper consumer CMOS image sensors (Section 5.2.3).

6.6.1.2 Depth, Colour and Size (DCS) algorithm

For implementation of the depth, colour and size algorithm, *H1* (Figure 5.9) in the DCSA segmentation technique was set to approximately the mid-point of the height of the pyrethrum (approximately 0.15 m) and *H2* was set 0.2 m higher than the pyrethrum (approximately 0.5 m). The components in the colour and depth image with heights below *H1* and above *H2* were separated into the deleted image while the components from heights between *H1* and *H2* were put in the retained image. A BST was applied to the deleted image and if any of the component total pixel size was above a user defined threshold $minWeed_{\text{thresh}}$, the component was deemed a weed. Plants smaller that small pyrethrum but between *H1* and *H2* are deleted in the DCSA.

6.6.1.3 Depth, Colour, Size and Spatial position (DCSS) algorithm

The depth, colour, size and spatial algorithm built up on the implementation of Section 6.6.1.2, and added spatial positioning of the components in the retained image. A component identified as pyrethrum by the DCS process in the retained image is then positionally compared to the centre of the image (a similar analysis as SP) and if the component was not positioned centrally it was deemed a weed.

6.6.1.4 LBP and Depth (LBPD)

The LBP had the best overall performance of methods in terms of accuracy and processing speed. The LBP had the second highest accuracy (55% compared to GLRLM at 63% Table 6.1) yet only required one pass over the image for rotational invariance compared to two passes for the GLRLM (0° and 90°). The C code implementation of the real-time rotationally invariant LBP algorithm developed in this thesis is given in Appendix F. The output of the LBP (or GLCM or GLRLM), on its own, was not satisfactory in terms of plant identification accuracy, and in order to enhance the LBP's effectiveness, average height and variance of the heights associated with the area of the image within the sliding window were added to the extracted LBP features. Average height and variance were chosen as the combination of these two features highlighted the height and evenness in height of pyrethrum plants. The features were then assessed by ROC curves to determine which were the most effective.

The three most effective features (depth variance, edge and flat) were applied to the classification techniques used in Section 6.3 and compared to the same training and data set of images as the initial evaluation of texture features. The results shown in Table 6.6 show the LBPD had a positive identification rate of 90% opposed to the original LBP at 55% (Table 6.2), highlighting the benefits of combining colour and depth data.

Table 6.6: LBPD classification results on pyrethrum.

Feature extraction technique	**SVM accuracy**	**MLP accuracy**	**k-NN accuracy**	**Naive Bayes accuracy**	**Training set size**	**Test set size**
LBPD	89.6%	90.1%	85.6%	90.1%	202	202

6.6.1.5 Effect of LBPD window size on images from real-world situations

A factor that arose while attempting to replicate the results of Table 6.6 on other pyrethrum images was the size of the window being classified. The window size was important as the images were multi-plant images and the DCSA segmented the occluding plant. However, some of the segmented components were too small to supply a consistently repeatable classification. This was similar to the findings in the sugarcane application of the system (Section 6.5.5.2).

In order for the texture operator to supply a consistent repeatable result it needed a minimum size sliding window, and from trial and error, a sliding window size of 32 × 32 pixels was identified. To determine the minimum window size, window sizes of 128 × 128, 64 × 64, 32 × 32, 16 × 16 and 8 × 8 were applied to sample images. Sliding window sizes were not identified as an area of concern in the literature review (Chapter 2) as the majority of research was conducted on still images (high resolution) and images with only one plant type in them making up the entire image window.

6.6.2 Evaluation of the developed techniques for feature extraction and classification in pyrethrum

The evaluation of the techniques developed for pyrethrum has been undertaken in two ways. Firstly to evaluate the hit rate, miss rate and false triggers on a pixel-by-pixel basis. Secondly, to evaluate the hit rate, miss rate and false triggers in relation to real-world identification of the total pyrethrum plant, and the possible damage to the pyrethrum by 'overspray'[2]. A false trigger typically involves a complete non-target plant being misclassified as a target plant; whereas

---

[2] Overspray is where the pyrethrum is misclassified at the start of the plant, and end of the plant, and therefore would be sprayed with herbicide causing damage to the pyrethrum.

overspray describes parts of a correctly classified plant containing a misclassified portion or portions.

Evaluations on a pixel-by-pixel basis supply results that define the identification accuracy of the algorithm as a proportion of the total of all pyrethrum plant pixels evaluated. Therefore a hit rate result of 95% means that 95% of the overall pixels of the pyrethrum plant evaluated were identified, not that 95% of individual pyrethrum plants were identified. The results from the pixel-by-pixel evaluation were used to rank the algorithms for use in the second evaluation process. The second process evaluated the hit rate of individual pyrethrum plants and what proportion of the plant was misclassified, causing overspray.

6.6.2.1 Results for pixel-by-pixel analysis

To undertake the pixel-by-pixel analysis an Automated Evaluation Application (AEA) was developed to automate the evaluation of the DCS, DCSS and LBPD techniques. The AEA compared a technique's classification results on an image (pixel-by-pixel) against a 'ground-truthed' image to validate accuracies, recording the results.

6.6.2.1.1 Real-time implementation effect on accuracy rates

Commercial spot spray herbicide delivery technology cannot spray weeds on the left or the right of the pyrethrum plant (with reference to the direction of travel of the spot spray system) without spraying the pyrethrum plant due to the minimum spray pattern width of the recommended spray nozzles. The commercial spray technology does have finer resolution in the direction of travel as this is controlled by the on/off electronic signals. Therefore, in order to enhance computation speed, the algorithms do not complete the classification of the pixels in the row (perpendicular to the direction of travel) once the row has been determined to contain pyrethrum, and this means the hit rate and miss rates do not always add up to 100%.

6.6.2.1.2 Ground-truthing

The images were ground-truthed by developing a second application. The second application required the operator to circle the weeds in the image with a mouse. On a separate colour image (initially blank), the co-ordinate positions inside the circled areas were filled automatically with the value 255 in the blue channel, to create a mask. Figure 6.6 shows the colour image with the white boundaries drawn around the weeds by a mouse and the subsequent mask of blue areas in the image.

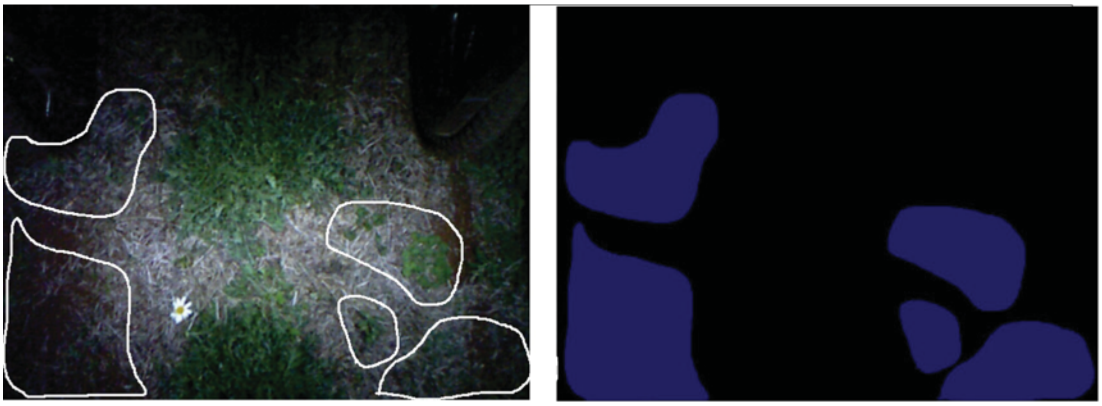

Figure 6.6: Images showing the ground truthing mask for the evaluation software. The image on the left shows the weeds circled by a mouse in white. The image on the right is the mask with the circled areas of the first image highlighted in blue.

6.6.2.1.3 Automated Evaluation Application (AEA) operation

The AEA initiated by accessing the video streams of the depth, colour and mask images, synchronising them so that each frame coincided to the same space on the ground. The AEA performed the analysis of the technique under scrutiny on the depth and colour images frame by frame. Wherever the technique identified a pyrethrum pixel, the AEA checked the mask image and if the same area in the mask was blue, then it was a misclassification. Likewise, everywhere the technique identified a weed pixel, the mask image was checked, and a blue pixel in the mask signified a correct identification. Results were collected from a sequence of video footage of 500 frames with the results from the classified pixels attributed to the following categories:

- *Pyrethrum pixels.* Plant pixels that were classified as pyrethrum in both the depth and colour test images and the ground-truthed image.
- *Missed pyrethrum pixels.* Plant pixels in the depth and colour test images classified as weed and appearing in the ground-truthed image as not weed (i.e. pyrethrum).
- *False positives.* Plant pixels in the depth and colour test images that were classified as pyrethrum but appearing in the ground-truthed image as weed.
- *Weed pixels.* Plant pixels in the depth and colour test images identified as weed and appearing in the ground-truthed image as weed.

6.6.2.1.4 Weed growth effects on results

The results have been segregated into two categories based on the weed growth range. The two ranges are: complete weed coverage (out-of-control weeds Figure 6.7) and intermittent coverage (in-control weeds Figure 6.8). When the weeds are out-of-control, they cover the ground around the pyrethrum completely making it difficult to visually determine where the pyrethrum is, or is not. A threshold on the percentage coverage of the ROI by plant material was used to identify the out-of-control condition in an image. Trial and error determined the threshold level at 90%. Therefore, out-of-control weeds are where plant material covers 90% and above of the ROI and in-control weeds are where there was less than 90% of the ROI covered by plant material.

Figure 6.7: Image showing out-of-control weeds in pyrethrum.

Figure 6.8: Image showing in-control weeds in pyrethrum.

6.6.2.1.5 Analysis of results for out-of-control weeds

Tables 6.7 and 6.8 contain the pyrethrum hit rate, pyrethrum miss rate and the false trigger rate in percentages relative to the total amount of pyrethrum pixels in the test images.

Table 6.7: Pixel identification classification results with respect to the total number of pyrethrum pixels for out-of-control weeds.

Out-of-control weeds			
Feature extraction method	**Miss rate**	**False trigger rate**	**Hit rate**
DCS	0%	8%	47%
DCSS	2%	8%	84%
LBPD	0%	9%	52%

Table 6.8: Pixel identification classification results with respect to the total number of pyrethrum pixels for in-control weeds.

In-control weeds			
Feature extraction method	**Miss rate**	**False trigger rate**	**Hit rate**
DCS	1%	3%	93%
DCSS	1%	1%	98%
LBPD	1%	2%	92%

From the results in Table 6.7, the response of the algorithms in areas where the weeds were out-of control was significantly worse than where the weeds were in-control (Table 6.8). Table 6.7 for out-of-control weed shows the range of correctly identified pyrethrum pixels was between 47% and 82% and the incorrectly identified weed pixels was between 8% and 9%. The aim of the spot sprayer in the field is to spray the weeds when they are in the in-control growth stage, and not the out-of-control stage, as it is too difficult to spray the weed without getting overspray on the crop. The results for the out-of-control area of the field show that the identification algorithms do not function satisfactorily in this field condition. Therefore, the following results analysis are determined from the in-control results.

6.6.2.1.6 Analysis of results for in-control weeds

Table 6.8 shows that the DCSS algorithm has the highest correct hit rate at 98% and the lowest miss rate at 1%. The LBPD algorithm appears to have the lowest performance of the three developed algorithms with a hit rate of 92% and a miss rate of 2%. The DCS algorithm performs slightly better than the LBPD with a 93% hit rate and a 3% miss rate. The DCSS and LBPD algorithms were advanced for further investigation in the second evaluation process. The DCS algorithm was a sub component in the DCSS algorithm.

6.6.2.2 Real-world pyrethrum identification results

The analysis methods have been developed on the hypothesis that **"if the pyrethrum pixels are identified, the remaining pixels are weed pixels"**. As introduced in Section 6.6.2 above, the second evaluation on the two algorithms determined the hit rate, miss rate and false trigger rate with respect to pyrethrum. Additionally, this second evaluation determined the pyrethrum plant's exposure to overspray from the spot spray process.

6.6.2.2.1 Real-world accuracy rate

The 510 frames of data from the in-control weed data set were visually inspected with the results shown in Table 6.9. Table 6.9 shows there were 78 pyrethrum plants of which 77 were correctly identified (98.7% hit rate and a 1.3% miss rate) in both algorithms and no misclassified weeds (0% false trigger rate). Visual inspection found that the missed pyrethrum plant had a low height component that did not meet the height criteria in either the DCSS or the LBPD algorithms.

Table 6.9: Pyrethrum accuracy.

Feature extraction	**Number of pyrethrum plants**	**Number of identified pyrethrum plants**	**Number of misclassified pyrethrum plants**	**Number of misclassified weeds**
DCSS	78	77	1	0
LBPD	78	77	1	0

6.6.2.2.2 Real-world overspray results

The number of plants in Table 6.10 with theoretical overspray was 10 (12.8% of data) for both algorithms. The amount of overspray was manually determined by comparing the amount of the individual pyrethrum plant with incorrect pixel classification to the overall size of that individual pyrethrum plant. Overspray occurred at the start (lead-in) or end (lead-out) of a plant and visual inspection of the plants in question showed that lead-in and lead-outs of the plants had low height. Therefore, the height criteria for pyrethrum were not met in the DCSS and LBPD and this was the cause of the misclassification. Table 6.10 shows average overspray in the DCSS and LBPD were similar at 9.5% and 10% respectively for the 10 plants but the variation in the range of the overspray was different with the DCSS being 6% and the LBPD being 11%. The average overspray relative to the whole dataset is 1.22% (9.5% × 12.8%) for the DCSS technique and 1.28% (10% × 12.8%) for the LBPD.

Table 6.10: Pyrethrum overspray evaluation.

Feature extraction	**Number of plants with overspray**	**Percentage overspray maximum**	**Percentage overspray minimum**	**Percentage overspray average**
DCSS	10	12%	6%	9.5%
LBPD	10	16%	5%	10%

6.6.3 Weed discrimination in pyrethrum – discussion and conclusions

6.6.3.1 Benefit of depth and colour data

The DCSA segmentation technique was able to adequately separate the occluding plants into individual components and the classification of these components then became a task for the feature extraction and classification techniques. At the outset, depth appeared to be a promising unique feature for pyrethrum identification, based on visual inspection of the data, and this was supported by the results of the LBPD in Table 6.6. Table 6.6 included depth as a feature, compared to the results of the original LBP in Table 6.2. These results (90% and 55% respectively), demonstrate an improvement of 35% was as a result of adding depth data to the algorithm.

6.6.3.2 Overspray error caused by sliding window

A source of error was at the lead-in, and lead-out, of the pyrethrum plant. The error was found at the changeover point, where the image transitions from ground to pyrethrum or pyrethrum to ground. At these transition points, the depth of the pyrethrum was varying from low to high, or high to low. The DCSS was able to detect this transition at the pixel row but the LBPD could not. The LBPD failed due to the sliding window analysis. The size of the sliding window fixed the resolution of position as the plant/weed transition would occur somewhere within the window. The lower resolution of positional accuracy produced the lower detection rate of the LBPD at 93%, compared to the DCSS of 98%, and in the range of overspray of the LBPD at 11%, compared to 6% for the DCSS.

6.6.3.3 Sliding window size

The sliding window required a minimum window size of 32 × 32 pixels, filled with pyrethrum for the window to be consistently identified by texture (LBP). Figure 6.9 shows a 32 × 32 pixel window bordered by white pixels inside a red ellipse. The window bordered by white is also a part of a pyrethrum plant that has a size of four, 32 × 32 pixel windows. Figure 6.10 shows a large pyrethrum plant with a size of 34, 32 × 32 pixel windows.

Figure 6.9: Image with four LBPD sliding windows of identification (highlighted green squares on the pyrethrum plant in the red ellipse). The green block bordered by white inside the red ellipse demonstrates the the size of a 32 × 32 pixel window.

Figure 6.10: Image with 34 LBPD sliding windows of identification (highlighted green squares on the pyrethrum plant in the red ellipse).

6.6.3.4 Best overall performance

The highest performing algorithm was the DCSS algorithm in terms of correctly identified pixels (98%) and of overspray which is displayed in Table 6.9. These combined analysis results show that for the evaluated pyrethrum[3] the performance was low on the out-of-control weed area, although this was less than 1% of the total field area. However, on the in-control weed growth areas, which was greater than 99% of the field area, the results showed a false positive rate of 1 in 78 plants (1.3%) and overspray of 1.2% of the total pyrethrum plant area. Therefore, the DCSS algorithm was determined to be successful at detecting weeds in pyrethrum on speedlings, grown in 0.65 m rows, at a growth stage of between 0.20 m and 0.45 m in height, and a minimum diameter of 0.15 m.

---

[3] Grown as 'speedlings' (planted seedlings) on the DRFSpeedlings site, which was the trial and test field provided by Botanical Resources Australia.

6.7 Summary of feature extraction and classification research

Chapter 6 has:

1. Evaluated existing texture extraction techniques (GLCM, GLRLM and LBP) with respect to sugarcane and pyrethrum.
2. Evaluated existing classification methods (MLP, SVM, k-NN and Naive Bayes) with respect to sugarcane and pyrethrum.
3. Developed custom feature extraction and classification techniques for commercial cropping field trials for sugarcane and pyrethrum.
4. Evaluated custom feature extraction and classification techniques for commercial cropping field trials for sugarcane and pyrethrum.

The evaluations found that:

1. Variation in plants was highlighted as a problem as no two plants were the same due to: the way the plants present themselves to the camera and damage due to die-back or potentially pest damage, weather damage, growth stage, nutrient and moisture availability.
2. The evaluation of the existing texture techniques and classifiers (Section 6.3) showed they could discriminate guinea grass and pyrethrum at between 49% and 65% accuracy. However, this was not adequate for real-time, real-world spot spraying and custom techniques were developed to improve upon the accuaracy rates.
3. An evaluation method (Section 6.4.1) was determined based on the standard criteria used in the literature review (Chapter 2), which was the hit and miss rate. This was augmented with the false trigger rate, because when

spot spraying weeds in an in-crop situation, a false trigger equals a dead crop plant. The ideal rates were identified as 100% hit rate, 0% miss rate and 0% false trigger rate.

4. The techniques developed in this research were shown to be successful for real-time, real-world plant identification (the in-depth real-time use of the algorithms is discussed further in Chapter 7). The developed algorithms can be adjusted to vary their hit rate and false trigger rates. In practice, the adjustment is targeted for a particular spray application (e.g. increasing hit rate at the expense of additional false triggers) and will vary depending on field conditions, weed infestation and tolerance for crop losses. The OTC had a guinea grass hit rate of 87% and a false trigger rate of 3.9%. Four techniques (Spatial, DCS, DCSS and LBPD) were developed for pyrethrum. DCSS was the highest performing algorithm with an identification rate on pyrethrum of 98% with an overspray of 1.2%.

5. The minimum component size required for identification was found to be a source of error in both the sugarcane algorithm and the pyrethrum algorithms.

- **Part software and part hardware**. Real-time systems can be implemented in hardware (e.g. Field Programmable Gate Arrays (FPGAs) and Digital Signal Processors (DSP)) for reasons of cost and performance, usually combined with microprocessors for control and communication.

Real-time machine vision weed identification systems reviewed in Chapter 2 (e.g. Wang et al. (2007), Berge et al. (2012) and Gerhards & Chrisensen (2003)), used microprocessor based hardware for analysis and real-time sequential computing methods. This was expected, due to the complexity of the algorithms being coded and the relative difficulty of programming an FPGA compared with a microprocessor. Multi-core processor integration in consumer electronics is a recent development which has been available since 2005 (Intel 2012).

7.2.3 Real-time systems operation

Berry (1989) defined computing system operation in three ways, using a level of interaction of the system with the environment as a basis:

1. **Transformational systems** are systems where the inputs to the system are supplied at the start of the process and a result is received at completion from the system, i.e. interaction with the system is at the start and end only.
2. **Interactive systems** are continually interacting with the environment around them. However, this interaction is at the system's pace, independent of changes in the environment.
3. **Reactive systems** are systems that are continually interacting with the environment around them at the environment's pace.

7.2.4 Real-time computational deadline terminology

Terminology for deadlines in a real-time system are listed below, provided by Bernat et al. (2001):

- Hard. A hard deadline is where a deadline cannot be missed because the consequences are great.
- Firm. A firm deadline is one where a task can miss a deadline but the result is useless.
- Soft. A soft deadline is where the the system can tolerate some deadlines missed and the result is still useful.

The component tasks for a real-time spot spray system defined in Section 4.1 and reproduced in Figure 7.1 are image acquisition, segmentation, feature extraction, classification and decision/action. Real-time sequential computing describes computational tasks being processed one after the other.

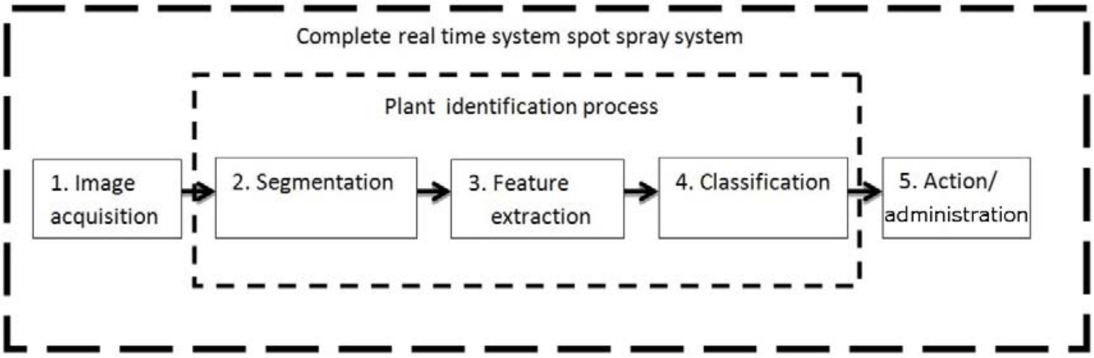

Figure 7.1: Reprint of block diagram of the spot spraying system based on real-time machine vision.

7.2.5 Real-time computation definition with respect to weed spot spraying

The term real-time refers to the application of computing system definitions. In most cases, real-time systems fall into the reactive computing definition category (Halbwachs 1993) outlined in Section 7.2.3 above. Definition 1 (Figure 7.2) combines the definitions of real-time and reactive computing to define weed spot spraying as a real-time reactive system.

Figure 7.2: *Definition 1* Real-time reactive machine vision weed spot spray system.

A real-time machine vision weed spot spray system must perform its actions whilst passing over crop rows and weeds (i.e. interacting with the environment) at the groundspeed being traveled by the agricultural vehicle on which the system is mounted (i.e. the environment's pace).

7.3 Real-time computing considerations for spot spraying

7.3.1 Object Identification Redundancy (OIR) between frames

Object Identification Redundancy (OIR) between frames refers to the minimum number of times the same position on the ground will have image analysis applied. OIR is required in no-till situations where the stubble may be concealing weeds from the image sensor or where larger plants are masking smaller plants.

Applying image analysis to the same position on the ground taken from different angles as the system moves over the ground improves the system's ability to view the weed satisfactorily from the occluding stubble. Figures 7.3(a) to 7.3(c) demonstrate occluding stubble in wheat stubble which is a principal broadacre crop in Australia.

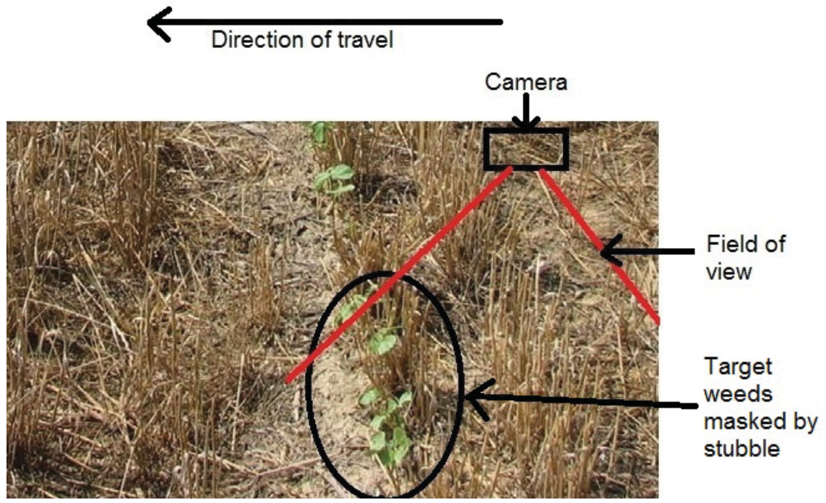

(a) Camera positioned right of weeds, occluded by wheat stubble.

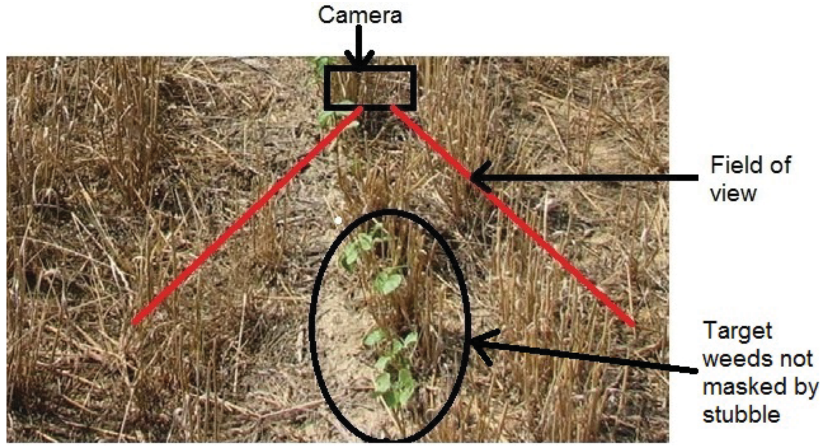

(b) Camera positioned above weeds and not occluded by wheat stubble.

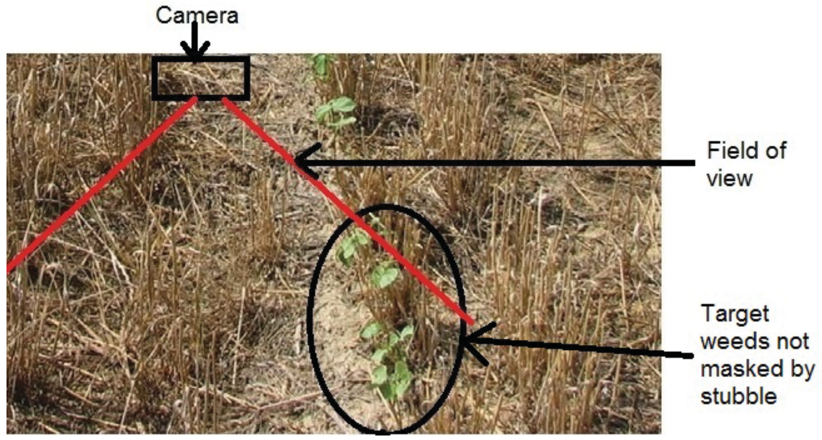

(c) Camera positioned left of weeds, not occluded by wheat stubble.

Figure 7.3: Sequence of images showing occlusion of weeds in wheat stubble and illustrating the need for OIR. The sequence starts with the camera positioned before the weeds and ends after the weeds. Occlusion is found in Figure 7.3(a) but the analysis would still identify the weeds in the following images. (original photograph by Northern Graingrowers).

In Figure 7.3(a), the camera is behind the weed position and the weed is masked by the stubble. In Figures 7.3(b) and 7.3(c), the camera is directly above and in front of the weed such that the view of the weed is not occluded in these images. For a fixed camera height, the number of frames for OIR is dependent upon stubble height and density in fallow situations and crop height and density for in crop situations.

7.3.2 Availability of computation time

The maximum groundspeed of the machine vision weed spot spray system is determined by the OIR of frames and the frame rate; and frame rate is determined by the computation time available. The following formula 7.1, relates minimum frame rate to groundspeed. The Formula 7.1 determines the interval between frames $t_{\max}$ therefore the frame rate $FR_{\min}$ is $1/t_{\max}$.

$$FR_{\min} = \frac{1}{t_{\max}} = \frac{n \times s}{d} \tag{7.1}$$

where:

$t_{\max}$ = interval between frames.
s = groundspeed in m/s.
n = minimum number of OIR frames.
d = length of the analysed ROI on the ground in metres.

For example, in pyrethrum, typical values of $n = 3$, $d = 0.4$ and $s = 4.2$ m/s were used for analysis at groundspeeds of 15 km/h to illustrate the computational time availability. From formula 7.1, $t_{\max} = 0.032$ s and $FR_{\min} = 31$ fps.

The maximum time available for processing in a sequential computing system (systems used in the literature review Chapter 2) is the time between frames, $t_{\max} = 0.032s$.

7.3.2.1 Consequences of computational overrun

The consequences of computational overrun are missing incoming frames, reduced reliability of OIR (as an individual frame that is missed may be the only frame which a weed is satisfactorily viewed) and poor synchronisation of the spray nozzle, i.e. spray pattern misses detected weed. The frequency of missed frames will be determined by the length of computational overrun.

Computational overrun is demonstrated in a timing diagram displayed in Figure 7.4. The top trace of the timing diagram in Figure 7.4 shows image frames acquired sequentially and identified by a separate frame number (based on an equipment frame rate of 30 fps which equates to a time interval of 33 ms between frames) and the acquisition time highlighted by the black, dotted line. The bottom trace displays the number of the frame as the frame is loaded into the analysis system, based on 45 ms processing time with the analysis acquisition time highlighted by the red line. The analysis of 45 ms is not adequate to keep up with the image acquisition of 33 ms and therefore, the system is forced to skip every fourth image.

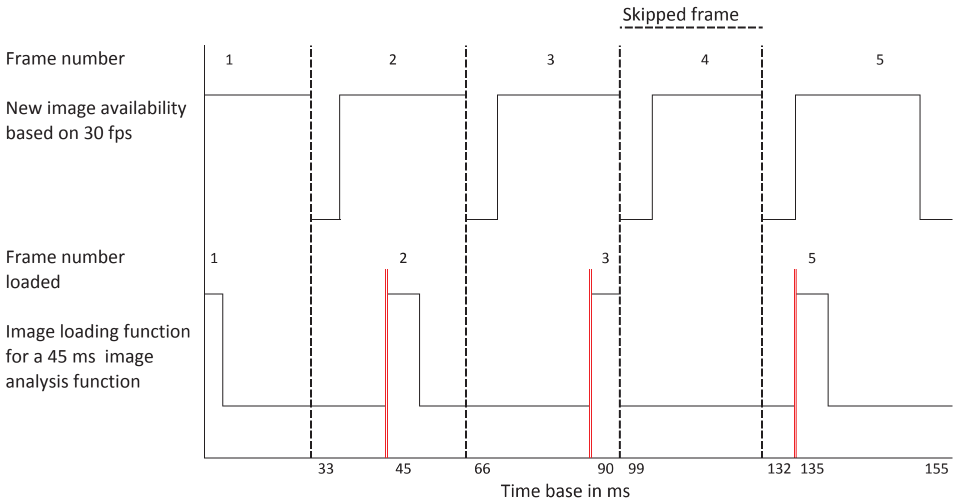

Figure 7.4: Timing digram showing an image analysis system missing an input frame of data due to the image analysis overrunning the image capture interval.

7.3.2.2 Computation time limitations

Processing tasks that are acted upon by outside stimulus (e.g. communication with the external spray system) may operate on CPU 'interrupts'. Interrupts cause the CPU to postpone the current task and undertake what is required by the interrupt, before returning to the postponed task. This approach may cause the interrupted task to push out completion time causing computational overrun. Possible interrupts need to be taken into account in the analysis time to meet 'hard' or 'firm' deadlines. Therefore, in the 15 km/h example (Section 7.3.2), there is a computation time of only 32 ms available to complete all tasks including interrupts.

7.3.3 Pipeline-based real-time systems

Hardware-based systems, sometimes referred to as embedded systems, typically use a logic device such as an FPGA for image analysis and may include an additional microprocessor for the supervisory tasks. A machine vision spot spray system was developed and commercialised by Rees Equipment Pty Ltd (Kinmont et al. 1999) (reproduced in Appendix E) using a logic device and microprocessor incorporating traditional logic pipelining (Figure 7.5) of the image analysis functions.

In a typical logic pipeline (Figure 7.5), data enters at logic process 1 and at each clock pulse the processed data is moved through to the next logic process, until the last logic process where the processed data becomes the output of the logic array. New data is input to process 1 of the logic array at every clock pulse but there is a delay associated with each logic process before the final result of the data is output at process N. The logic array could have any number of processes and the data would be delayed for more clock pulses but each clock cycle would return a result. The only stipulation is that each process should have an execution time of less than the clock pulse period.

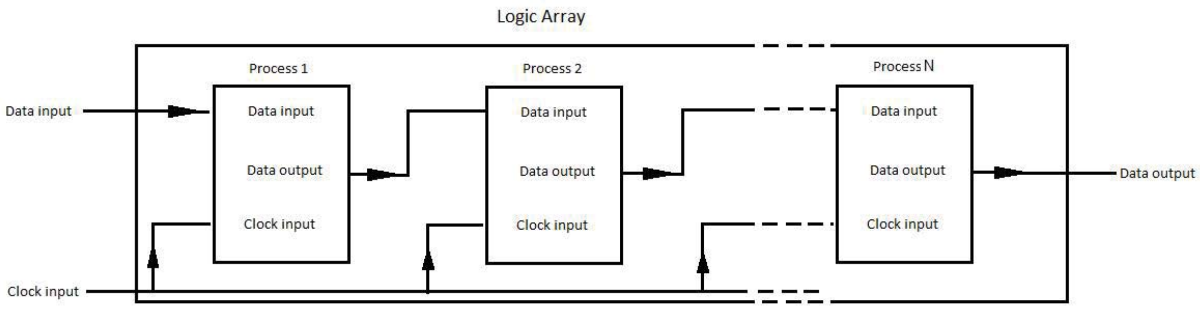

Figure 7.5: Block logic array illustrating data flow through an N-process hardware logic pipeline controlled by synchronous clocking.

A benefit of the logic pipeline in the spot spray system of Kinmont et al. (1999) was that additional processing could be accommodated, but with a consequent lengthening of the delay between the image input at process 1 and output at the final process. The delay between input and output can be accounted for when implementing nozzle synchronisation.

7.3.4 Nozzle synchronisation

A correlation between the weed on the ground and the same weed's position in the image must be made to accurately spray herbicide onto the weed. The correlation can be made by synchronisation the time delay between when the image was acquired and when the weed should appear below the nozzle. The time delay is a function of the weed's position in the image, the groundspeed of the system and the distance between the centre of the image and the nozzle. The time delay is affected when the time that the image analysis acts upon the image (red line of the bottom trace in Figure 7.4), varies with respect to the actual image acquisition time (black dotted lines of the top trace Figure 7.4), creating a variable time delay in activating the nozzle. Therefore the time from image acquisition (relates to the physical position on the ground) to the action of turning on the spray nozzle at the end of the analysis process is varying. The variation in the time delay means that herbicide may be delivered before or after the weed and precise timing of herbicide application onto weeds cannot be guaranteed.

7.3.5 Conclusion

There is a need for superior real-time computing approaches. Hence, a custom approach has been developed as set out in Section 7.5.1 below based on multi-core and parallel processing techniques. These latter techniques are first reviewed in Section 7.4 following.

7.4 Review of single and multi-core processing

Consumer electronics have undergone significant advancements in the last several years with the development of cheap multi-core processors. Multiple CPU computers have been available prior to the introduction of consumer multi-core processors (e.g the Burroughs D825 modular data processing system (Anderson et al. 1962)) but were not utilised by the mainstream, as the computers were expensive and the implementation of tasks was complex, requiring developers with specific skills (Campbell & Miller 2010). However, along with consumer multi-core processor advancement, new development tools implementing parallel processing techniques have been provided by companies (e.g Intel® and Microsoft®) which have made programming for parallel processors available to the mainstream software developer.

Development tools can have variations in terminology. For example, Intel® uses 'parallel building blocks' to describe parallel functions and Microsoft® use 'parallel patterns'. For consistency, this thesis uses Microsoft® terminology.

7.4.1 Typical single core programming methods

7.4.1.1 Sequential processing and concurrency

Typical single core processing methods are 'sequential' and 'concurrent'. 'Sequential' processing was the first reported processing technique used in 1945 by John von Neumann, and refers to a process being executed and the system waiting until the process is finished before moving on to the next process (Akhter & Roberts 2006). 'Concurrency' allows more than one process to operate at the same time and is obtained through multi-tasking and multi-threading capabilities which were developed in the 1960s to make use of the computing system's resources in the most efficient manner (Akhter & Roberts 2006). Concurrency overcomes the limitation of sequential functions and makes use of the time lost waiting for user input. However, execution times cannot be guaranteed because the operating system does not know when an interuption may occur (e.g when a callback[1] might be triggered).

7.4.1.2 Multi-tasking and multi-threading

Multi-tasking is achieved by the operating system allowing each process (task) to execute in small time slices and changing from one process to the next based on a priority schedule administered by the operating system (Intel 2003). This gives

---

[1] A callback is a function (or pointer to a function) that is passed as an argument to another function, which is expected to execute the argument at a time when triggered. The execution can be immediate as in sequential systems or any time later as determined by the operating system in concurrent systems (Laksberg 2012).

the illusion of running more than one process concurrently. A further evolution was added with multi-threading which allows the processes to split into smaller functions called threads which can then be scheduled to operate in a similar way to multi-tasking; or can be left dormant and only called on when some other action has occurred (Intel 2003).

As an example, if a camera is sending data to the computer via USB, the USB driver will trigger a callback when the computer has received the frame of data. The operating system must then fit the workload from the callback into the scheduling program to act on the data.

7.4.1.3 Consequences of single core programming methods for spot spraying

Concurrent operation is a primary method employed in commercial operating systems and is seamlessly integrated into the development tools such as the Microsoft® Visual Studio suite of products. This approach gives control of the concurrent process scheduling to the operating system. The drawback of concurrent processes in a real-time system such as weed spot spraying, where the time constraints are *hard*, is that execution times cannot be guaranteed. This uncertainty can add to the processing time and cause the overall processing time to exceed the allowable time for the computation of a result.

7.4.2 Parallel processing

In parallel processing complete tasks, individual functions or low level instructions are allocated by the operating system and typically executed asynchronously on individual processor cores or divided amongst multiple cores to execute at the same time. The principal benefit of parallel processing is the speed up in execution time by spreading the processing load amongst multiple cores (Campbell & Miller

2010). Akhter & Roberts (2006) state that although the overview of parallel processing may sound similar to concurrency, the terms are not interchangeable. When a number of threads or tasks are running in parallel, they are all running simultaneously on different hardware processors. When a number of threads or tasks are running concurrently they are all running on the one hardware processor with their own allocated time slice. Akhter & Roberts (2006) state 'In order to have parallelism, you must have concurrency exploiting multiple hardware resources'.

The current philosophy for programming parallel processes is similar to the concurrent operation, previously outlined in Section 7.4.1, and the drawback is the same, which is that the system cannot guarantee execution times. This may not be a problem if the overall speed increase in the system is so great, from the use of the multi-cores that even in the worst case, the processing time is still within the required time frame. However, the speed increase is not linearly related to the number of cores used, and is the focus of Amdahl's Law, outlined in the next section.

7.4.2.1 Amdahl's Law

Amhahl's Law determines the relative speed improvement to a software program by parallelising its operation (Amdahl 1967). Not all portions of a software program are able to be parallelised (i.e. must be left as sequential) and the possible speed improvement of a program's operation using parallel processing is limited by this sequential portion. Amdahl's Law calculates the possible speed-up in processing as:

$$speed_{\mathrm{improvement}}(f,n) = \frac{1}{(1-f)\frac{f}{n}} \tag{7.2}$$

where:

n = the number of processor cores.

f = the amount of the program that can be parallelised.

Figure 7.6 is a graph of Amdahl's Law from Equation 7.2 with $f = 75\%$ (typical value from computing texts further discussed in Section 7.6.1.2) and n varying from 0 to 20. Figure 7.6 shows the speed-up from four cores is approximately 2.25 times, nine cores is a speed-up of three times and 20 cores is a speed-up of 3.5 times emphasising a non linear speed improvement. The non-linear improvement provides diminishing returns for speed increase by the addition of extra processing cores.

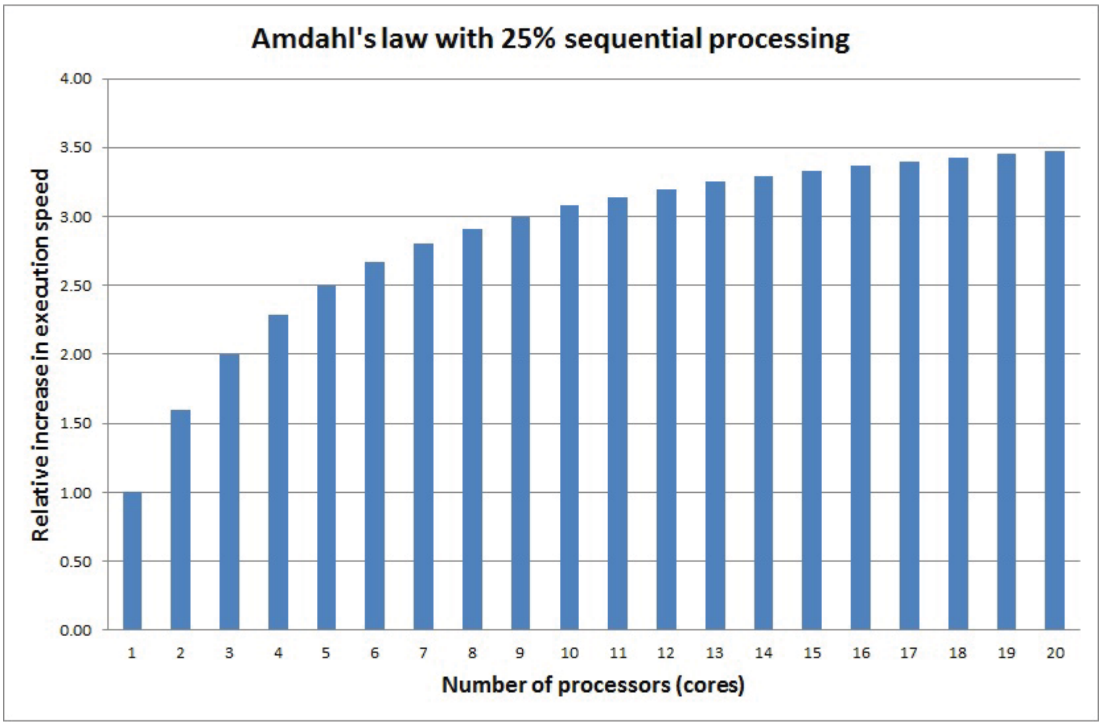

Figure 7.6: Graphical representation of the relative execution speed improvement determined by Amdahl's law with 25% sequential processing over 20 cores.

7.4.2.2 'Parallel' patterns in the Microsoft® development platform

Campbell & Miller (2010) outlines the Microsoft® development platform parallel patterns implementations as:

1. Parallel loops. A calculation is performed on the contents of a data group

with no dependencies[2] from the preceding or following item of data.

2. Parallel tasks. The operating system runs tasks as separable asynchronous tasks.
3. Parallel aggregation. Parallel aggregation provides similar outputs to parallel loops and is used when there are dependencies.
4. Futures. The outputs of some operations are used as inputs into other operations and the order in which the operations are constrained. The operations may or may not be able to run in parallel depending on the data dependency.
5. Dynamic task parallelism. Tasks are dynamically added as the computation proceeds such as in database sorting.
6. Pipelines. A pipeline is where the output of one task (stage) is fed in as the input of another task.

The six patterns above are used in asynchronous parallel processing implementations where the operating system allocates the functions.

7.5 Novel Synchronised Pipeline Processing (SPP) technique

The novel Synchronised Pipeline Processing (SPP) technique developed in this research incorporates traditional parallel computing patterns and hardware pipelining of the image analysis algorithms to extend the processing time available with a linear improvement. The speed-up supplied by the synchronous pipelining

---

[2] A dependency refers to the relationship between software functions. Function (A) is dependent on function (B) if function (A) requires input data from the output of function (B).

method is a speed-up in the input frame rate, not an overall speed-up in execution time of the software functions. This combination has not (to the author's knowledge) been published, and an opinion on patentability (reproduced in Appendix D) suggests that the developed technique has not been used in the context of real-time image processing.

7.5.1 Modified pipeline used in the SPP

The principal pattern that the SPP technique modifies is the pipeline. Typical asynchronous pipeline operations and the improved synchronous pipeline operation with SPP are discussed in the following subsections.

7.5.1.1 Asynchronous pipeline operation

Figure 7.7 illustrates an asynchronous parallel process pipeline function. The pipeline process flow is similar to hardware flow (Figure 7.5) but the system is asynchronous. The memory buffers at the end of each stage in Figure 7.7 need to be large enough to hold multiple frame's worth of data, as each stage does not pass on the results synchronously. A stage may have to wait for the results from the previous stage; or if a stage is slow it must store data from the previous stage. At some point the output of the system may need to slow down to accommodate the slowest function or else data will be lost. If the pipeline is running multiple frames the tasks will be reallocated for each frame. This system cannot guarantee execution time which is not satisfactory in a real-time system.

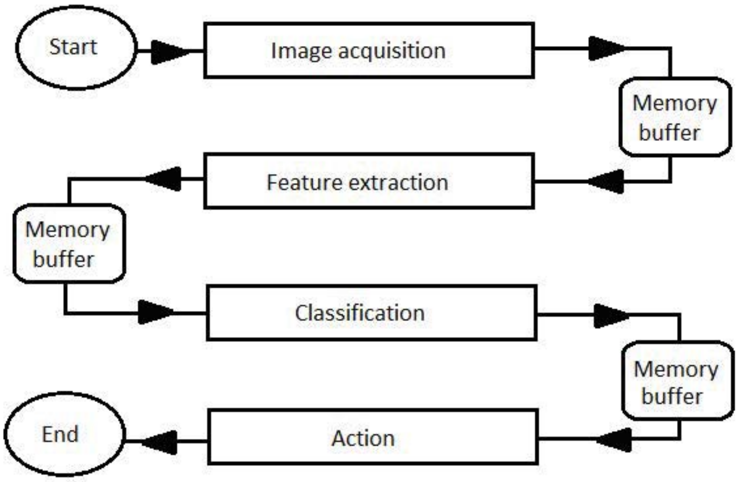

Figure 7.7: Asynchronous parallel pipeline process applied to image analysis (adapted from (Campbell & Miller 2010)).

7.5.1.2 SPP concept and application to spot spraying

The SPP can guarantee execution time and overcomes the non-linear improvement associated with typical asynchronous parallel patterns. The SPP achieves this by allocating fixed tasks to the individual cores which operate continuously in an endless loop until the spot spray application is complete (e.g. end of the agricultural field). The fixed tasks are synchronised to each other by the incoming image frame from the camera or a timer synchronised to the incoming image frame.

Figure 7.8 demonstrates the SPP technique as applied to the real-time processing requirements of spot spraying. All of the analysis functions associated with each stage of the processing are allocated to an individual core and synchronised by the input frame number N. Cores P0 and P5 are operated asynchronously as they use the asynchronous functions of the operating system (image acquisition and computer inputs and outputs). Cores P0 and P5 operate with a synchronous hardware timer so that they cannot overrun the allocated time frame and skip frames. Cores P2, P3 and P4 operate sequentially and are synchronised by the input image. In cores P2, P3 and P4 executions are timed so that in a worst case they still will not overrun the allocated time frame determined by the frame rate.

In this example (Figure 7.8), cores P2 and P3 are operating in parallel with the feature extraction function split between the two cores. Alternatively it could be two different feature extraction functions.

Figure 7.8 is one example of the implementation, however the technique can be applied to multicore CPUs and multicore GPUs with many more or fewer cores. The allocation of the analysis process to the cores is based on the dependencies of the analysis processes.

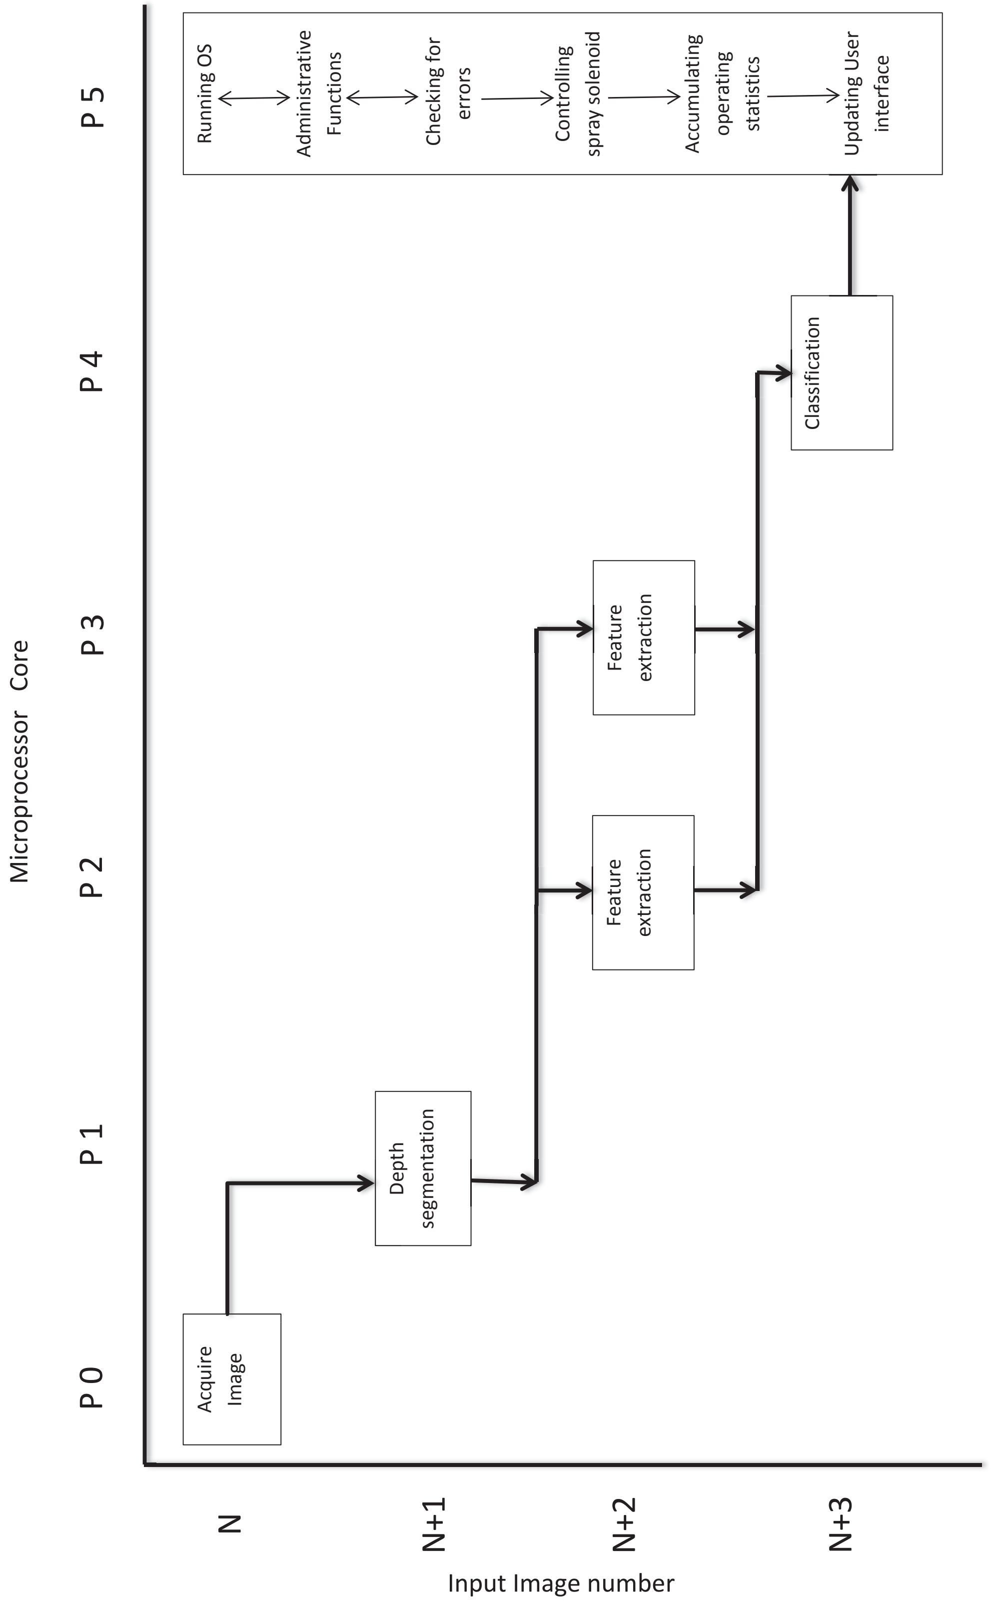

Figure 7.8: Flow diagram of the Synchronised Pipeline Processing (SPP) technique, indicating the principal processing tasks in its application to spot spraying.

7.5.1.3 Timing diagram for SPP operation

Figure 7.9 is a timing diagram of the SPP technique with consecutive input frame numbers set forth horizontally across the diagram. The processing stages of the pipeline are set forth vertically downward beside the diagram and the designator P corresponds to the core allocation in Figure 7.8. It can be seen that there is a delay from when input frame number 1 enters the pipeline at P1 and when the output for frame number 1 is actioned in the last processing stage of the pipeline at P5. It can also be seen that no frames are lost, and that there are now five input frames worth of acquisition and processing time available. Additional processing time can be obtained by using a processor with a higher core count in the same way.

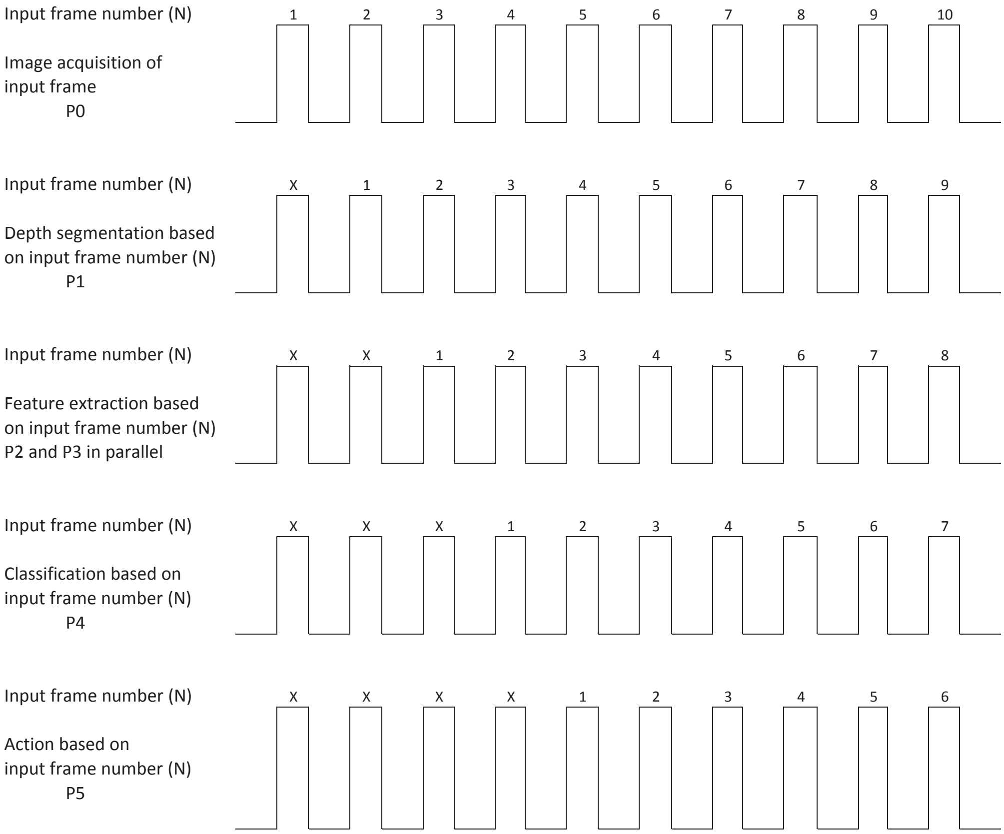

Figure 7.9: Timing diagram of the novel SPP technique corresponding to Figure 7.8. X denotes an unknown state of the system.

The system has the following flow through the stages relative to the time period:

- **P0**- The images are acquired (frame number N=1) and pre-processing is performed with the images being made available to stage 2 at the end of the time period.
- **P1**- The pre-processed images of frame number N=1 are taken by the depth analysis system in P1 of the pipeline, analysed and the result made ready for P2 and P3. P0 is repeated with frame number N=2.
- **P2 and P3**- The result from P1 analysis on frame number N=1 is processed in P2 and P3 which include two feature extraction algorithms running in parallel on different cores with results made available for P4. Extra feature extraction algorithms could be run at this stage if implemented on a processor with a higher core count. P0 is repeated with frame number N=3 and P1 is repeated with frame number N=2.
- **P4**- The results of P3 analysis on frame number N=1 are classified and the result passed to administration. P0 is repeated with frame number N=4, P1 is repeated with frame number N=3 and P2 and P3 are repeated with frame number N=2.
- **P5**- Output for frame number N=1 is actioned and P0 is repeated with frame number N=5, P1 is repeated with frame number N=4, P2 and P3 are repeated with frame number N=3 and P4 is repeated with frame number N=2.

The above sequence repeats continuously until the overall process is ended.

7.5.1.3.1 Summary of SPP operation

The similarity of the SPP to the flow of a hardware logic device (Figure 7.5) is demonstrated by the flow of results from processing stage to processing stage, clocked by either a hardware timer, set to be synchronised with the input images frame rate, or the new input image itself. The processing in stage P1 is part sequential and part asynchronous with the overall possible processing time including the asynchronous function being less than the input image frame rate period. The processing within stages P2 to P4 are 100% sequential and execution time is fixed. Administration and action functions P5 are asynchronous. The functions executing within administration do have their scheduling manipulated programatically to influence the operating system's timing allocations by changing the thread priority settings (Windows SetThreadPriority function set to THREAD PRIORITY TIME CRITICAL (MSDN 2015)), placing more emphasis on real-time tasks.

7.6 Evaluation of alternative processing configurations

This section evaluates the processing times for:

- sequential single core processing;
- asynchronous parallel processing as used by Microsoft®; and
- the novel SPP technique.

The frame rate available from each image analysis application in sugarcane and pyrethrum is determined because frame rate is a principal factor in determining groundspeed (set out in Equation 7.1).

7.6.1 Evaluation method

The method used to compare the processing techniques is to determine the execution times for each Image Analysis Task (IAT) in a sequential single CPU core and apply the times to asynchronous parallel processing techniques and the the SPP.

7.6.1.1 Execution timing data collection

A possible source of execution timing error is the operating system interrupting the IAT to address other functions of the computer or callbacks. The operating system interruptions can distort the execution timing results and one method of addressing the operating system interruptions is to collect execution timing results from numerous frames and determine an average (Persa et al. 2000). The total time taken for the weed spot spray function on a single CPU is the addition of the IAT times of all the modules in milliseconds.

7.6.1.2 Asynchronous parallel processing timing setup

The speed increase factor for Microsoft® asynchronous parallel processing was calculated to be 2.25 using Amdahl's Law of parallelisation with 75% parallelised (Figure 7.6) based on a quad core CPU operating at 2.8GHz.

The portion of code that can be parallelised is difficult to estimate as each loop within a function needs to be evaluated. The amount of parallelisable code in each of the software functions developed for spot spraying in this research would be expected to vary due to the dependency of data within the function. For example a function for binarising and filtering an image has few dependencies in the data and could be estimated at approximately 90%, whereas the DCSA is heavility dependant upon previous data and could be estimated at approximately 10% paralellisable.

To fully determine the portion of parallelisable code in the sugarcane and pyrethrum algorithms, would require a complete re-write of all software functions to implement the parallel patterns. Therefore, for the purpose of evaluation a conservative value of 75% paralellisable code was choosen for all functions combined. The total asynchronous parallel processing execution time was determined by dividing the single CPU time by 2.25.

7.6.1.3 SPP timing setup

The distribution of the analysis modules to individual cores similar to the allocation shown in Figure 7.8 was required to determine the SPP execution time. The largest core operation time will be the fastest execution time usable for the SPP method to ensure that all processes, in all cores, can be carried out. Dependencies of the modules are taken into account when the modules are allocated to individual cores.

7.6.1.4 Computer setup

The computer used was a 2.8 GHz, Intel$^{®}$ I7 2640M with 8GB RAM. The setup of the computer when collecting the execution timing results was:

- no other programs were open;
- all external communications were disabled.
- the thread running the program had a priority setting of ‘high priority’ (the highest setting possible);
- the program was set to core 2; and
- the execution timing of the analysis modules is determined by the ‘Windows’ function QueryPerformanceCounter(), which returns the current value of the high-resolution performance counter in 100 ns increments.

7.6.1.5 Frames per second calculation

The frames per second results were determined by using the total execution times for single core and asynchronous parallel processing methods and the fastest time for the SPP method in milliseconds and dividing one second by this value.

7.6.2 Results and discussion – sugarcane

The execution timing data for sugarcane was acquired from 1,996 consecutive frames of the video data set from 10/10/2012, field 13-A (Table 3.1) which was used for the feature extraction and classification results of Section 6.5.5. There are seven analysis modules in the weed spot spray function for sugarcane and these are listed in Table 7.1. Table 7.2 displays the module times taken for the sequential processing of each analysis module. Table 7.3 shows the allocation of modules to processing cores.

Table 7.1: Analysis modules for sugarcane.

Analysis Module	**Process number**
Image acquisition	1
Segment and filter colour image	2
Pre-process depth image	3
DCSA	4
Delete all leaves above *h2*	5
Combine the retained image and colour image	6
Classify weed by tracking algorithm	7

Table 7.2: Execution times of analysis modules for sugarcane.

Timing statistic	**Analysis module**						
	1	2	3	4	5	6	7
execution time (ms)	4	6.71	8.73	24.49	1.28	6.58	10.61

Table 7.3: Allocation of the analysis modules to the individual cores for sugarcane SPP analysis.

Processing core	**Core 1**	**Core 2**	**Core 3**	**Core 4**
Analysis module allocation	(1) (2) (3)	half of (4)	half of (4)	(5) (6) (7)
Running time per core per frame (ms)	19.44	12.24	12.24	18.84

The functionality of the DCSA is in two halves, i.e. half the analysis is done in core 2 and at the next frame the partial result from core 2 is completed in core 3 before being passed onto core 4 at the following input frame. Figure 7.10 highlights the improvement in frame rate of the parallel methods over sequential.The overall execution time and implied maximum frame rate for each method with fps shown in Figure 7.10 was as follows:

- The single CPU was 62.40 ms (corresponding to 16 fps).
- The asynchronous parallel processing execution time was 28 ms (36 fps).
- The SPP execution time was 19.44 ms (51 fps).

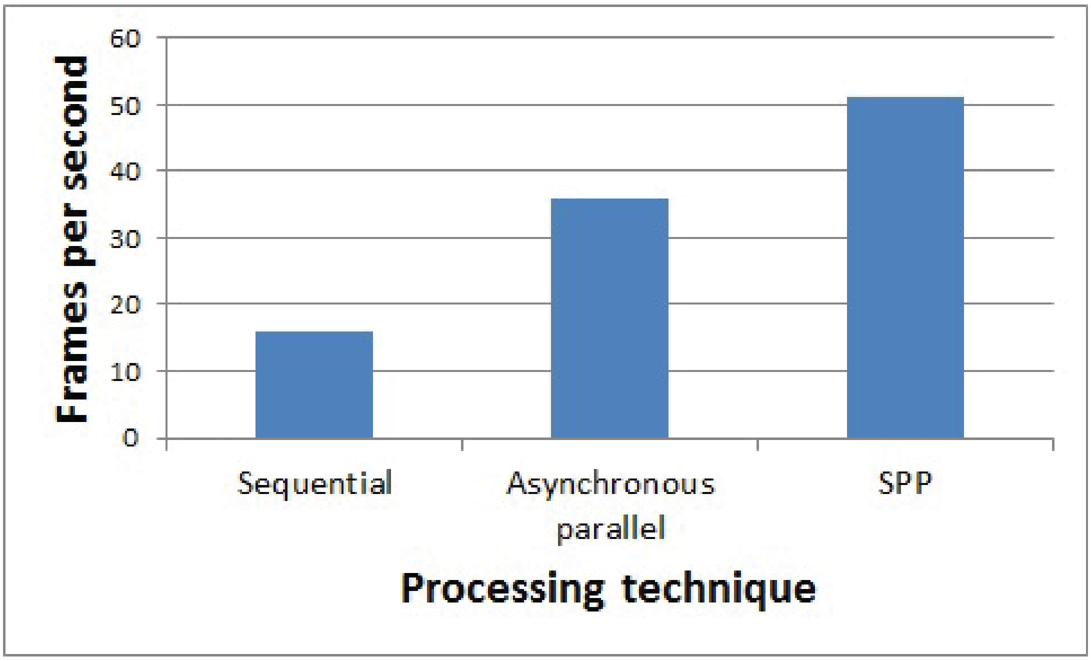

Figure 7.10: Overall improvement in frame rate of processing techniques for sugarcane.

By comparing the execution times for each processing technique and the frame rate results of Figure 7.10, the SPP requires 31% of the processing time of sequential processing, which is a speed-up in frame rate of 3.2 times. The same analysis for asynchronous parallel processing shows 44% of the processing time compared to a sequential processing method (a speed-up of 2.25 times).

7.6.3 Results and discussion – pyrethrum

The execution timing data for pyrethrum was acquired from 500 sequential frames of data taken from the DRF-Speedlings site on the 11-04-13 (Table 3.3).

The weed spot spray software function for pyrethrum has 10 analysis modules, listed in Table 7.4. As with the timing data for sugarcane, Table 7.4 displays the analysis modules and Table 7.5 displays the module execution times taken for the sequential processing of each analysis module. Table 7.6 displays the allocation of the analysis modules to the individual cores (with dependencies taken into account). Figure 7.11 displays the frame rates of the processing methods.

Table 7.4: Analysis modules for pyrethrum.

Analysis Module	**Process number**
Image acquisition	1
Depth calibration	2
DCSA	3
Copy images	4
Spatial segment	5
Binarised image segmentation	6
Hull guidance	7
LBP classification	8
Depth classification	9
Combined depth and spatial classification	10

Table 7.5: Execution times of analysis modules.

Timing statistic	**Analysis module**									
	1	2	3	4	5	6	7	8	9	10
execution time (ms)	14.79	1.15	21.95	0.68	38.65	11.49	2.04	3.07	1.06	1.99

Table 7.6: Timing of the individual cores when the analysis modules have been distributed. Times are in milliseconds.

Processing core	**Core 1**	**Core 2**	**Core 3**	**Core 4**
Analysis module allocation	(1) (2) (6)	(3) (4)	half (5) (7) (8)	half (5) (9) (10)
Running time per core per frame (ms)	27.43	22.63	24.44	22.38

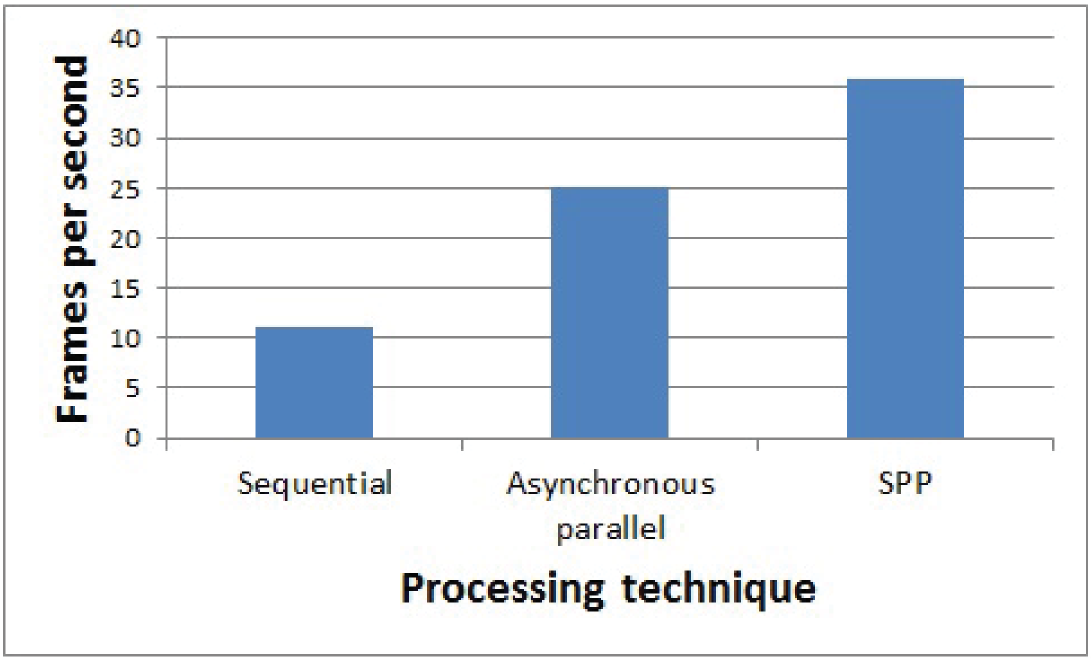

Figure 7.11: Overall improvement in fps of compared techniques for pyrethrum.

Table 7.6 shows analysis modules one, two and six are allocated to core one; modules three and four are allocated to core two; the functionality of module five in two halves, i.e. half the analysis is done in core three and at the next input frame the partial result from core three is completed in core four. From Table 7.6, the largest time in the core operation is 27.43 ms in core one. Hence core one will determine the fastest analysis time usable for the SPP method in pyrethrum.

The overall execution time and implied maximum frame rate (fps) for each method shown in Figure7.11 was as follows:

- Single CPU is 90.76 ms (corresponding to 11 fps).
- The asynchronous parallel processing was 40.3 ms (25 fps).
- The SPP was 27.43 ms (36 fps).

From the execution times and the frames per second of each method listed above, the SPP is 29% of the processing time compared to sequential processing, which is a speed-up in frame rate of 3.35 times. The asynchronous parallel processing execution time provided 44% of the processing time compared to a sequential processing method (speedup of 2.25 times).

7.7 Discussion and significance of synchronous pipelining to spot spraying

7.7.1 Groundspeed improvement

A primary factor in a real-time, real-world spot spray system is the input frame rate as it dictates the groundspeed of the spot sprayer as discussed in Section 7.3.2. The 'speed-up' of the input frame rate supplied by the synchronous parallel pipelining method compared to sequential processing and asynchronous parallel processing is summarised in Table 7.7. Although the overall processing times of sugarcane (62.4 ms) and pyrethrum (90.79 ms) varied by 45%, the SPP had a similar improvement over sequential and asynchronous parallel processing. The similar improvement is because the operation of the techniques did not change between crops and therefore the speedups remained relative.

Table 7.7: Summary of the 'speed-up' of input frame rate created by the SPP method compared to sequential processing and asynchronous parallel processing.

'Speed-up'	**Sugarcane**	**Pyrethrum**
over sequential processing	3.21	3.35
over parallel asynchronous processing	1.41	1.49

The benefits of increased processing capacity of the SPP for spot spraying can be used a number of ways:

1. increase in groundspeed of the spot spray system; or if not required then
2. increase in sensor data and complexity of algorithms to identify weeds and crop on-the-go and apply herbicide to the targeted plant; or if not required then
3. increase in robustness of result by operating secondary identification algorithms to check original algorithm's result; and/or
4. use a cheaper computer with lower specifications to achieve the practical maximum groundspeed.

To highlight the advantage in groundspeed of the SPP in the field for sugarcane and pyrethrum:

- the data from the sugarcane timing example (Section 7.6.2) can be used. Maximum groundspeeds for the different processing techniques can be determined by using the frame rate in Figure 7.10 and a practical value of 10 for CF_{thresh} (Section 6.5.1). Therefore, sequential processing would have a maximum groundspeed of 5.7 km/h, asynchronous parallel processing techniques would have a maximum groundspeed of 13 km/h and the novel SPP technique would have a maximum groundspeed of 18.5 km/h. Practical, commercial speeds in sugarcane are less than 8 km/h so the extra processing capacity could be used as outlined in items 1 to 4 above.

- the pyrethrum frame rates with an OIR of 3 and $d = 0.4$ (Section 7.3.2) sequential processing would have a maximum groundspeed of 5.25 km/h, asynchronous parallel processing techniques would have a maximum groundspeed of 12 km/h and the novel SPP technique would have a maximum groundspeed of 17.2 km/h. As with sugarcane practical, commercial speeds in pyrethrum are less than 8 km/h so the extra processing capacity could be used as outlined in items 1 to 4 above.

7.7.2 Nozzle offset for commercial weed spot spraying

The computational time taken (total pipeline analysis time) between the input of the frame and the output of a result (77.76 ms in the sugarcane example Section 7.6.2 and 109 ms in the pyrethrum example Section 7.6.3) at the end of the pipeline in the SPP can be compensated for by positioning the spray nozzle a distance behind the camera that allows for the total pipeline analysis time and maximum groundspeed. The distance required to move the nozzles can be determined by multiplying the maximum groundspeed in m/s by the total processing and solenoid activation time in seconds. Therefore for a maximum groundspeed of 8 km/h (2.2 m/s) and solenoid activation time of 0.010 s (10 ms), the mounting distance offset for in sugarcane will be 2.2 X (0.077+0.010) = 0.19m (190 mm) and pyrethrum will be 2.2 X (0.109+0.010) = 0.26 m (260 mm).

7.7.3 Further improvements

The improvement in frame rate (or the capacity for additional analysis) will be greater with a higher core count CPU compared to sequential or asynchronous parallel processing. This is because the sequential processing portion of the software will stay the same; the improvement in asynchronous processing will not be linear (Amdahl's Law) and will taper off; but the increase in improvement for the synchronous pipeline method will continue until core allocation of the analysis modules cannot be further achieved due to the dependency of modules or sub modules. The speed-up for the sequential pipeline will then follow Amdahl's Law as further improvements will be made by 'spreading the load' allocated to each core over more cores as in the asynchronous methods.

7.7.4 Conclusion

It has been demonstrated that the SPP can provide an advancement in the amount of processing available to identify the weeds from the crop in two different farming industries in Australia. In the sugar and pyrethrum industries where the groundspeed would not exceed 8 km/h, the SPP will provide the capacity to undertake additional feature extraction and classification analysis such as applying additional texture and shape extraction techniques. The additional analysis can be used to identify additional weeds and expand the system's capabilities of identification, or to improve the accuracy of the classification with more complex analysis operating in real-time. In industries that have a high groundspeed requirement, such as the broadacre industry in Australia which operate at up to 20 km/h, the SPP will allow analysis to distinguish crop from weeds at higher groundspeeds than have been available to date.

Chapter 8

Portability of the Depth Colour Segmentation Algorithm (DCSA)

8.1 Introduction

This chapter reports a preliminary evaluation of the portability of the Depth Colour Segmentation Algorithm (DCSA) to crops other than sugarcane and pyrethrum and also the operation of the DCSA within the Synchronised Parallel Pipeline (SPP). The crops used for the evaluation of portability were sorghum and mung beans, chosen because they are common crops in the north eastern farming areas of Australia and are commonly grown in minimum and no-till farming. Sorghum and mung beans also represent grass-like crop (sorghum) and broadleaf type crop (mung beans) which are the principal crop categories in the broadacre and row crop farming sectors in Australia.

Feature extraction and classification is not evaluated in this chapter as the present research has not set out to develop a generic technique that is portable to other crops. However, in principle at least, the DCSA should be generic and therefore portable to other crops. The same analysis method for the DCSA's level of