University of Southern Queensland

Faculty of Health, Engineering & Sciences

Precision weed detection via colour and depth data fusion in real-time for automatic spot spraying

A thesis submitted by

Steven Rees

in fulfillment of the requirements of

Doctor of Philosophy

2015

Abstract

Broadacre and row crop farming in Australia uses no-till and minimum-till farming systems which have led to the overuse of specific herbicides for weed control, causing resistance to those specific herbicides to build-up in weeds. Automatic weed spot spraying can help reduce resistance build-up by specifically detecting weeds for targeted control with alternative herbicides, hence breaking the resistance cycle. Existing commercial weed spot sprayers are only capable of distinguishing green from brown, i.e. plant material in a fallow situation, and image analysis research for weed discrimination typically was not developed for commercial on-farm conditions. The research in this thesis has developed a real-time, real-world machine vision spot spray system that can operate at groundspeeds up to 20 km/h and discriminate green from green (i.e. weed from crop) under commercial conditions for two very different crop types and size scales, specifically sugarcane (grass-like) and pyrethrum (broadleaf).

Occlusion of a weed leaf by another leaf or plant is a major impediment for real-world operation of a machine vision weed spot sprayer. A Depth Colour Segmentation Algorithm (DCSA) has been developed which combines depth data and colour image data to segment individual leaves from each other, based on pixel connectedness in height and colour, providing an accuracy when occluded of greater than 99%. The DCSA has a filtering capability that can reduce the amount of data requiring further analysis by an observed 83% for sugarcane and 53% for pyrethrum.

Existing feature extraction techniques have been evaluated in the thesis and have been shown to be unsatisfactory in discriminating weed from crop especially when the weed and crop are similar species. e.g. grass-like weed (guinea grass) from grass-like crop (sugarcane). Depth features were added to the extracted features of a local binary pattern function, improving the accuracy from 63% to 90% for pyrethrum identification, and showing that depth data combined with 2D data can improve the discrimination result. Additional real-world custom algorithms have been developed to achieve an identification accuracy of 87% (where 86% of the weed was occluded) with a 3.5% false positive rate for sugarcane. The Depth, Colour, Size and Spatial (DCSS) algorithm developed for pyrethrum achieved 98% accuracy for pyrethrum identification with a 1.2% false positive rate.

Real-time functionality has been obtained by the development of a Synchronised Parallel Processing (SPP) technique. The SPP technique maintains a high frame rate (which determines the maximum groundspeed) by assigning the workload in a permanently allocated pipeline synchronised by the incoming video image. Calculations for sugarcane and pyrethrum show that speeds up to 18.5 and 17.2 km/h respectively are achievable based on the algorithms developed and a higher core count CPU (six cores were used in the calculation) would achieve higher groundspeeds. The gains from the additional processing availability provided by SPP can be used to achieve a higher groundspeed, or undertake additional image analysis, if required.

It is concluded that the machine vision components developed in this thesis comprise a real-time, real-world machine vision spot sprayer that can operate at commercial groundspeeds up to 20 km/h and discriminate weed from crop.

Associated Publications

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.

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

Rees, S (2014) 'Methods, systems, and devices relating to real-time object identification' U.S. Provisional Application 61/988,541, preservation date May 5, 2014.

Rees, S (2014) 'Methods, Systems and Devices Relating to Real-Time Object Identification'. PCT Patent Application No. PCT/US15/29261 preservation date May 5, 2014.

Certification of Dissertation

I certify that the ideas, designs and experimental work, results, analyses and conclusions set out in this dissertation are entirely my own effort, except where otherwise indicated and acknowledged.

I further certify that the work is original and has not been previously submitted for assessment in any other course or institution, except where specifically stated.

STEVEN REES   W0023751

_____   _____

Signature of Candidate   Date

ENDORSEMENT

_____   _____

Signature of Supervisor   Date

_____   _____

Signature of Supervisor   Date

Acknowledgments

Firstly I would like to thank God for leading me to undertake this research and giving me the ability to do it.

Thank you to my family for encouraging me during this time especially for reminding me during the write-up phase that it would end at some point.

I appreciate the support from the NCEA, Botanical Resources Australia, Bundaberg Farms 'Fairymead', Ian Ladner, Steve Hanlon and my colleagues at the NCEA.

Finally, thank you to my supervisors Dr Cheryl McCarthy and Associate Professor Nigel Hancock for your guidance.

Steven Rees

University of Southern Queensland
24th July 2015

Contents

List of Figures

List of Tables

Chapter 1

Introduction and Overview

As farming practices in Australian agriculture have changed, so too have weed control methods. Callow et al. (2010) state that the increased uptake of conservation farming practices has seen an increase in herbicide usage for the control of weeds. The increase in chemical usage is causing selective breeding of tolerance characteristics (also known as 'resistance' in the farming industries) to the chemical groups[1] used for weed control, in successive generations of the weeds. Callow et al. (2010) propose that a strategy to break this tolerance is to use a herbicide from a different chemical group with a different killing action and/or mechanical cultivation.

Automatic spot spraying[2] of weeds can be used to achieve cost effective weed control with herbicides (Cropoptics Australia 2012) but current commercial technologies are limited to differentiating green from brown i.e. a weed in a fallow[3] situation.

---

[1]Herbicides are classified into groups from A to N, based upon their killing action (Queensland Department of Primary Industries & Fisheries 2005).

[2]The spot application of a herbicide to an individual or small patch of weed or crop.

[3]Land that is left unseeded with crop and weed free during a growing season (Encyclopedia Britannica 2013).

Machine vision has the potential to determine a plant's species in real-time based on the analysis of the plant's visual features, and this may allow a rotation of herbicides to be used in controlling weeds and so break the resistance cycle. There has been considerable research in individual species identification by machine vision, as cited by Slaughter et al. (2008) and Ji et al. (2009) in their reviews on machine vision weed identification. However, both reviews found that the machine vision weed identification research undertaken to date had been developed for controlled conditions, and is consequently not suitable for achieving a workable commercial solution with general application.

1.1   Research aim

The aim of the research presented in this thesis is to develop a 'real-time' precision in-field sensing system that can discriminate crop from weed for the purpose of automatic weed spot spraying and weed mapping[4] in a practical agricultural setting. The scope of this thesis does not include the physical spray application system.

1.2   Hypothesis

The hypothesis of this research is that the identification of plant leaves in three dimensional (3D) space will improve image segmentation and isolate individual plants from occlusions[5] for further more intensive identification analysis.

---

[4]The recording of the weed and its latitude and longitude so that it can be displayed on geographic information software for evaluation.

[5]An occlusion is where parts of plants and/or stubble overlap each other from the viewpoint of the camera.

1.3   Specific research objectives

The objectives of this research are as follows:

1. To develop an algorithm/s that incorporates 2D (colour) and depth data from video streams to achieve weed discrimination from crop in a real-time, real-world environment (Section 1.9) at commercially realistic ground-speeds.

2. To demonstrate that the addition of depth data to a suitable image analysis technique can achieve weed discrimination from crop in a commercially acceptable operational window, i.e. at a range of crop growth stages, in a real-time, real-world environment.

3. To evaluate the performance of the developed technique under a range of real-world environment conditions; in particular with respect to 3D space versus 2D, or depth, on their own.

4. To demonstrate that the system is adaptable to a range of crops under practical commercial conditions.

1.4   Farming industries targeted

Particular emphasis was placed on the sugarcane (*Saccharum officinarum*) (Appendix C.1) and pyrethrum (*Saccharum officinarum*) (Appendix C.2) farming industries with preliminary evaluation on sorghum (*Sorghum bicolor*) and mung beans (*Vigna radiata*). All crops are illustrated in Figure 1.1.

Grass and broadleaf crops are the principal crop categories in the broadacre and row crop farming sectors in Australian farming. Sugarcane and pyrethrum are representative of a grass-like crop and a broadleaf-like crop respectively as well as

(a) Sorghum at mid season growth stage.   (b) pyrethrum at early stage growth.

(c) Sugarcane mid-season growth stage.   (d) Mung bean at early season growth stage.

Figure 1.1: Example of crops evaluated.

being of very different physical scale. The techniques developed in this research must be able to operate in the real-world environment of conservation farming, which has varying field conditions e.g. soil type, stubble cover, crop height and weed type.

1.5   Background – Conservation farming practices

The motivation for the research reported in this thesis arises from the need to improve conservation farming practices.

Conservation farming, which can also be referred to as conservation agriculture, was first evaluated in the 1940s in the United States as a means of soil erosion control (Fergal 2010). Figure 1.2 highlights erosion during the 'dustbowl' conditions in the U.S. in the 1940s. There have been advances in machinery and herbicides during the 1960s, 1970s and the 1980s which have accelerated the uptake of this farming practice (Fergal 2010).

Figure 1.2: 'Dustbowl' in the U.S 1940s (Wilson 2013).

Conservation farming is defined by the The United Nations Food & Agricultural Organisation (2012) as a farming philosophy that strives for a sustainable farming environment as well as improved farming profitability. These goals are achieved through three main principles: minimum soil disturbance, permanent soil cover and crop rotations. The Department of Primary Industries NSW (2009) states that the significant benefits of conservation farming are:

- An improvement in controlling erosion from both wind and water.

- Increased water storage and retention. The stubble (crop residue) cover

decreases the formation of crust on the surface of the soil which aids in water infiltration. The stubble cover also produces a buffer for soil temperatures which decreases evaporation.

- Increased soil biological activity, organic matter and organisms.

- Improved soil structure due to less tilling[6].

- Increased nitrogen content over time.

Common conservation farming practices in the Australian broadacre, row crop, sugar and pyrethrum farming industries are labeled 'no-till' and 'minimum-till'. No-till farming is a strategy where there is no tilling of the soil. The seed is planted through the residue (stubble) on the soil surface. The stubble should be maintained at the highest possible density and all weed control is done with herbicides. This approach maximises the benefits of conservation farming but is critically dependent upon effective herbicide usage.

Minimum-tillage is where farmers may sometimes utilise tilling[7] for weed control, preparation of the seed bed and seeding. Minimum-till generally incorporates some weed control by herbicides to reduce the overall amount of tilling required.

1.6 Background – Herbicide delivery systems

Herbicide weed control is achieved by the use of a boomspray which sprays herbicide onto the weeds. Boomsprays can be used to spray both herbicides and nutrients on fallow, crop, and between the rows of crop in row crop farming. Typical boom sprayers in the Australian marketplace with their operating groundspeed ranges, are:

---

[6]Tilling is the disturbance of the soil.

[7]Weed control is achieved with tilling by disturbing the root zone of the soil and cutting the plant's roots to kill the plant.

- self propelled, 15-35 km/h (Figure 1.3);

- trailed, 5-25 km/h (Figure 1.4); and

- between row shielded sprayer, 0-10 km/h (Figure 1.5).

Figure 1.3: Agchem high clearance sprayer (Croplands Australia Ltd 2013).

Figure 1.4: Croplands trailing boom (Croplands Australia Ltd 2013).

Figure 1.5: Wylie row crop shielded sprayer (Wylie sprayers Inc 2013).

1.6.1   Herbicide resistance

One of the challenges of minimum and no-till farming practices is the shift in weed spectrum that has occurred due to herbicide resistant (hard to kill) weeds. The overuse of a chemical group, or too low an application rate, for weed control causes selective breeding of tolerance characteristics to that chemical group in successive generations (Callow et al. 2010).

Figure 1.6 is a diagrammatic representation of the resistance cycle. Figure 1.6A shows the weeds being sprayed with a chemical at a rate that a weed with natural resistance can tolerate. All the other weeds die but the naturally resistant weed sets seed, shown as a single plant in Figure 1.6B. The weeds grown from the seeds are sprayed again and there are more naturally resistant weeds as shown in Figure 1.6C and the seed set cycle repeats. Figure 1.6D shows that at the end of several cycles most of the weeds are naturally resistant.
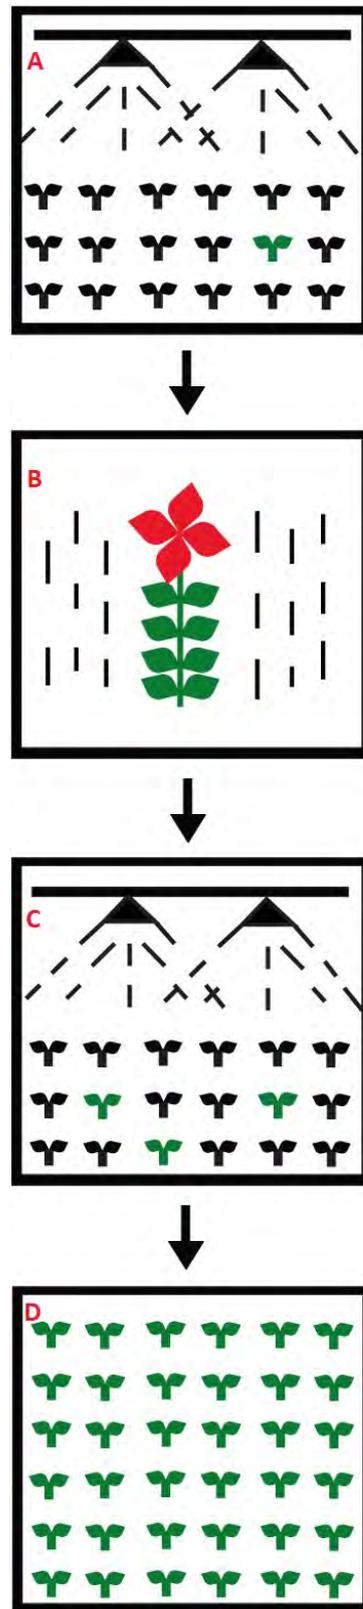
Figure 1.6: Herbicide resistance evolution flow chart, based on University of Minnesota (2008).

Known control methods (Callow et al. 2010) are either the application of a herbicide from a chemical group, with a different killing action, to that of the herbicide to which the weeds have become resistant, or mechanical disturbance by tilling. This is effective because changes of chemical group, when undertaking the spraying application in Figure 1.6C, will redirect the resistance cycle to Figure 1.6A, never arriving at Figure 1.6D, thereby permitting the reuse of the original herbicide.

One means to control weeds that is promoted by the Cotton Research Development Corporation (CRDC), the Grains Research Development Corporation (GRDC) and the Sugar industry is to implement an Integrated Weed Management strategy (IWM), as explained below in Section 1.6.2.

1.6.2 Integrated Weed Management (IWM)

Cotton CRC Australia (2010) define IWM as the use of numerous weed control tools and strategies in a long term plan to reduce the weed seed bank and control weed competition for crops. A weed may survive a single weed control approach but the chance of its survival is greatly reduced when faced with several different approaches. An IWM is summarised by a cycle of four categories (Australian Weed Management 2004) and shown in Figure 1.7.
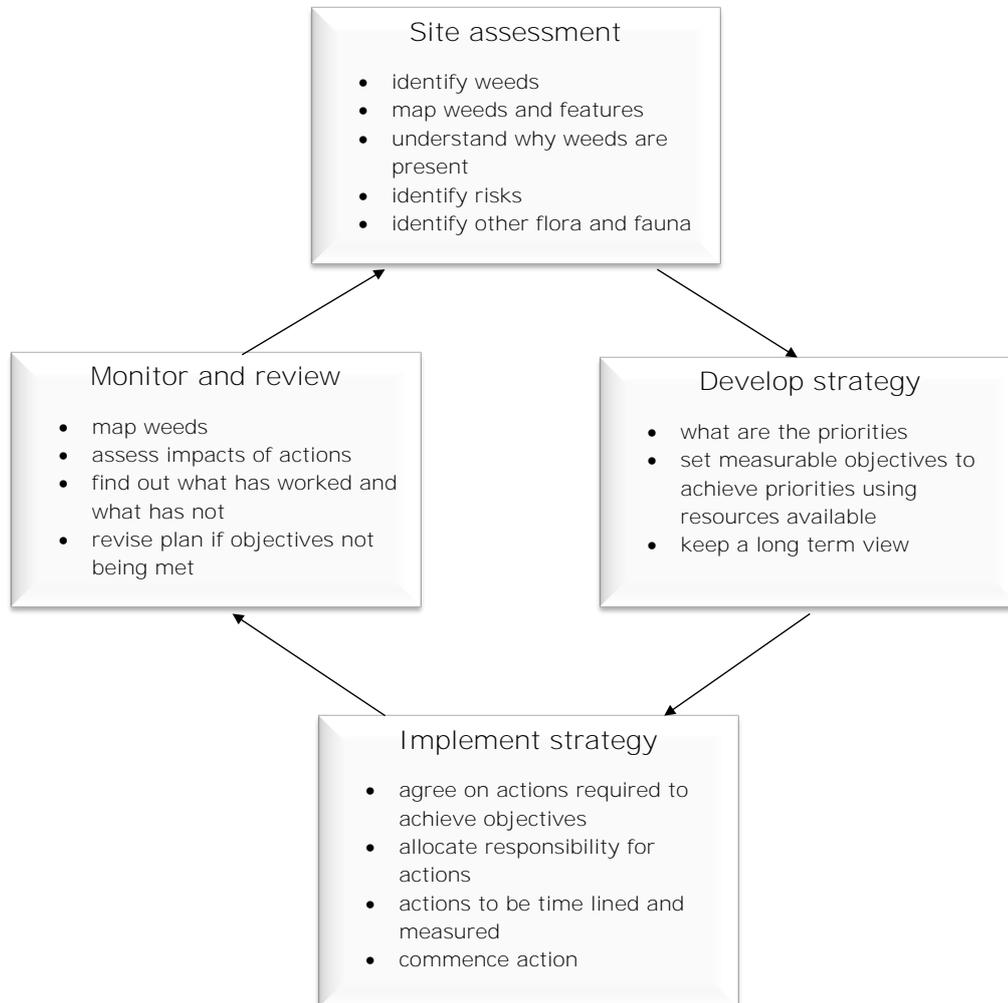
Figure 1.7: Weed management plan (adapted from Australian Weed Management (2004)).

Weed identification/spot spraying technology can become an extra tool in the IWM toolbox by: (i) providing an automated means of site assessment (weed identification); (ii) implementing a strategy of spot spraying with herbicides; and (iii) monitoring and reviewing by mapping the species.

1.7   Automated spot spraying

Spot spraying is the spot application of a herbicide to an individual, or small patch, of weed or crop. The process can be manual (triggered by a person with a backpack sprayer) or automated through a boom sprayer to allow the farmer to cover a large amount of farmland. Machine vision identification of weeds can be found in the literature as early as 1983 (Haggar et al. 1983) and commercial automated spot spraying has been available in Australia since 1984 (Cropoptics Australia 2012). Farmer testimonies indicate chemical savings of between 50% and 90% over the course of keeping a fallow weed free using an automated spot sprayer (Cropoptics Australia 2012).

Commercially available weed detection technologies (Figures 1.8 and 1.9) for automated spot spraying are limited to distinguishing plant from background (i.e. soil and/or stubble) and are outlined in Appendix B.

Figure 1.8: Boomspray equipped with Weedseeker technology (reproduced from Cropoptics Australia (2012)).

Figure 1.9: Boomspray equipped with Weed-It technology (reproduced from Croplands Australia Ltd (2013)).

Identifying the presence of vegetation allows the spot spraying of weeds in a fallow situation (Figure 1.10), i.e. plant from background, but not weeds in a crop (Figure 1.11), or individual plant species (plant from plant).

The algorithms developed in this research are able to identify weeds in a crop situation by overcoming the issues of occlusion, thereby not relying on weeds being on a soil or stubble background to be identified.

Figure 1.10: Fallowed paddock with wheat stubble and large weeds image taken at Felton, Qld Australia, July 2011.

Figure 1.11: Sugarcane field with grass weeds in the front of the rows image taken at 'Fairymead' farm, Bundaberg Australia, August 2012.

1.8   Machine vision for plant recognition

The literature review (Chapter 2) describes the methods and techniques that have been used to achieve weed from crop discrimination. The diversity of the literature shows that a large knowledge base exists for machine vision techniques including a combination of depth, spatial, spectral, shape and texture features.

A further requirement for a machine vision plant identification system, when incorporated into a automated spot sprayer, is real-time operation. Real-time operation requires the machine vision system to determine a correct result within a fixed amount of time (Lin & Burke 1992). Failure to achieve this can result in weeds being missed by the detection system.

1.9   Real-time computation and (agricultural) real-world conditions

Lin & Burke (1992) stated that a "real-time computer must produce a correct

result within a specified time" and West (2001) defines a real-time system as a system where "the results will be provided when they are needed". Hence a real-time, field deployable machine vision system for weed identification must be able to discriminate crop from weed as it progresses along the field at an effective working speed determined by the type of spraying being done, the crop and the ground conditions. Real-world conditions provide a myriad of variations: in stubble covers, plants at differing growth stages, and plants with different levels of health. For the purposes of this research, real-world conditions do not include rain, dew, high wind or high temperature, as the operator would not be spraying in these conditions (Primary Industries Standing Committee 82 2002).

In a review of autonomous robotic weed control systems, Slaughter et al. (2008) found that:

- the main obstacle holding back the commercial success of robotic weed control systems is the lack of robust weed detection;

- much of the research had been done under ideal conditions with correct identification results of between 65% and 95%, but that these results were not repeatable under field trials;

- the most common problem was occlusion of weeds from other plants; and

- variation in leaf appearance can vary greatly in the same species as a result of plant health, physical trauma, shadows and daytime conditions.

In addition Ji et al. (2009) and Slaughter et al. (2008) found that most research to date used static images whereas a real-time system must use a video stream.

1.9.1 Overcoming limitations in real-time weed identification – Occlusion

As noted above in Section 1.9, Slaughter et al. (2008) determined that the most significant source of error for identification of weeds based on imaging systems is occlusions. Occlusions create errors in segmentation as traditional segmentation techniques cannot find the edge between the overlapping leaves to separate them into individual components. Occlusions can change the perceived shape of the segmented plant component by making segmented components look larger or smaller (if the plant is occluded by stubble or other plants).

Research presented in this thesis demonstrates that the addition of depth information can allow automated weed discrimination to use the heights of plants, leaves and stubble to reconstruct individual leaves and/or remove unwanted leaves. This is demonstrated in the sequence of images in Figure 1.12. Typically, colour images of weeds (Figure 1.12(a)) are segmented into a binary image containing green pixels and background pixels (Figure 1.12(b)). Visual identification of individual plants leaves in the binary image is not possible when leaves overlap one another as the overlapping leaves combine in the same segmented component. However, with combined colour and depth segmentation (Figure 1.12(c)) overlapping leaves may be separated. As a result, the proportion of error in leaf shape determinations is reduced, and correspondingly the proportion of error in weed identification. Details of combined depth and colour segmentation are presented in Chapter 5.

1.9.2 Overcoming limitations in real-time weed identification – Real-time requirements

By definition, real-time systems must work within processing time constraints. Figure 1.13 shows a block diagram of a real-time machine vision system. Each of the functions in Figure 1.13 are required to happen within a specified time. To

(a) Colour image of sugarcane and guinea grass to be segmented. The red circle encloses an example of multiple overlapping leaves

(b) Binarised image of (*a*).

(c) New colour/ depth segmentation technique applied to image (*a*).

Figure 1.12: Image sequence highlighting occlusions and segmentation by colour and combined colour and depth.

address these constraints, a new processing technique for image analysis of weeds has been developed and evaluated in this thesis (Chapter 7). The new processing technique incorporates parallel computing methodologies in the Microsoft® software development platform, as well as hardware pipelining of the plant identification process in Figure 1.13, to maximise the processing time available.

The hardware pipelining strategies leverage off concepts developed for image analysis on logic devices (Field Programmable Gate Arrays (FPGAs) and Complex Programmable Logic Devices (CPLDs)) and can be implemented in microprocessors using the multi-core architecture available from CPU manufacturers (e.g. Intel®) as set out in Chapter 7.
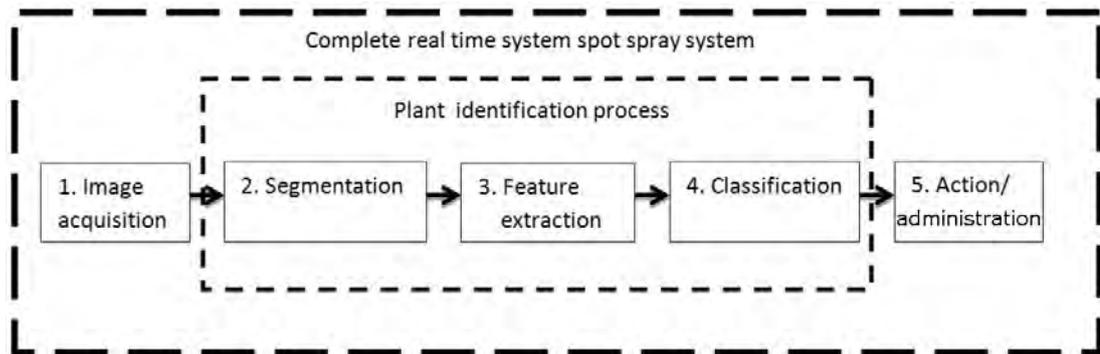
Figure 1.13: Real-time machine vision block diagram, displaying the complete spot spray system incorporated by the outer block and the plant identification incorporated by in the inner block.

1.10   Innovation and novelty

The novel research contributions reported in this thesis are as follows:

1. A colour and depth segmentation algorithm (Chapter 5).

2. A real-time implementation of a rotationally invariant Local Binary Pattern (Chapter 6 and Appendix F).

3. A real-time synchronous pipeline processing technique (Chapter 7).

4. The integration of the innovations above with traditional image feature extraction techniques and classifiers (Chapter 6).

The techniques of items 1 and 3 have been evaluated by Fisher, Adams, Kelly (FAK Patent Attorneys) for novelty and patentability (Appendix D). The outcome of the evaluation was that patent protection should be sought and the University of Southern Queensland has been actively pursuing this along with a commercialisation strategy and a commercial partner. Hence publication of the novel techniques in open literature has been limited. A provisional patent and PCT patent have been lodged and are itemised in the Publications section.

1.11 Dissertation chapter outline

The research undertaken for this thesis is aimed at developing a real-time, real-world automatic spot spray system to spray weed from crop. A complete real-time spot spray system comprises five main areas which are addressed in Chapters 4 to 6 with the real-time capability in Chapter 7 and algorithm portability in Chapter 8.

Chapter 2 is a literature review of current machine vision techniques that have been applied to weed identification. The chapter reviews segmentation of plant from background, identification based on shape, texture, spectral differences and 3D imagery. Strengths and weaknesses in the current research were determined from this review with limitations of processing speed, occlusion and inconsistent illumination being highlighted as issues that need to be addressed.

Chapter 3 is concerned with data collection. Data was collected from sugarcane crops at 'Fairymead' farm Bundaberg, Queensland Australia and pyrethrum crops by Botanical Resources Australia in Tasmania. Additional data to evaluate

portability of the algorithms was collected from sorghum and mung bean crops on the Darling Downs in Queensland Australia.

Chapter 4 outlines the methodology of a real-time machine vision system for weed identification in agriculture highlighting five steps (Figure 1.13). The chapter includes discussion of step 1, image acquisition and pre-processing. Pre-processing colour and depth images for a real-time system involves image remapping algorithms and procedures to ensure that an object appears in the same image co-ordinates in the colour images and the depth image. The remapping procedures discussed are for overlaying the depth and colour images accounting for the offset in the images due to the mounting distance between the cameras; and the time offset between the images due to timing differences in the cameras and the groundspeed of the system.

Chapter 5 addresses step 2 in the machine vision methodology outlined in Chapter 4, the segmentation of plant and background pixels. The chapter highlights the limited use of existing segmentation techniques, when applied to crop and weed segmentation in real-world conditions, and then discusses the novel Depth Colour Segmentation Algorithm (DCSA) developed in this thesis. The chapter describes the operation of the DCSA and the DCSA's filtering capabilities. The remainder of the chapter evaluates the application of the DCSA in sugarcane and pyrethrum, showcasing the DCSA's capabilities with respect to occlusion.

Chapter 6 considers the steps 3 (feature extraction) and 4 (classification) in the machine vision methodology outlined in Chapter 4. Existing (published) 2D feature extraction and classification techniques for weed-from-crop identification in real-world conditions are evaluated and found to be unsatisfactory. The chapter then outlines new techniques developed specifically for sugarcane and cotton in this thesis, which include colour and depth information. A methodology for evaluating the new feature extraction and classification techniques, to determine the most effective technique is developed and applied to the results of field trials in commercial sugarcane and pyrethrum cropping.

Chapter 7 details real-time processing and how processing speed limitations may be overcome. The limitations are outlined and expanded upon with a new technique developed that allows significant additional processing when compared to sequential processing, or typical parallel processing techniques. A performance assessment is undertaken on sugarcane and pyrethrum field data.

Chapter 8 addresses the application of the DCSA technique to two additional crops, sorghum and mung beans, highlighting the DCSA's real-time capability to effectively isolate plant components when occluded. This chapter illustrates the DCSA's portability to other crops in the Australian no-till farming sector, and its ability to operate in real-time, real-world conditions.

Chapter 9 sets out the conclusions and the possibilities for further work. The conclusion addresses the objectives individually to show how each objective has been successfully met. Further work looks at what more needs to be done to enhance practical commercial spot spraying.

Appendices

The seven appendices comprise:

A − Glossary of terms.

B − Commercial spot sprayers.

C − Introduction to sugarcane and pyrethrum farming practices.

D − Patent attorney opinion.

E − Hardware system patent.

F − Real-time LBP implementation.

G − Side shift hitch operation for row guidance.

Chapter 2

Literature review for weed identification with machine vision

2.1 Introduction

Machine vision can potentially identify weeds in real-time by using spectral information as well as spatial information where spatial information refers to a plant's position with respect to the crop row. The research reviewed in this chapter demonstrates that machine vision techniques have been used successfully to determine plant identification under particular operational conditions. These operational conditions are of limited use in the commercial world of farming due to real-world conditions (Slaughter et al. 2008). For example Tian et al. (1997) developed a system that operated at a specific (cotyledon) growth stage, however climatic conditions and external work pressures mean that it is not always possible for farmers to get onto the fields at a specific growth stage.

Machine vision systems may comprise one or more imaging devices that provide spectral and spatial data. Common 2D imaging cameras use a Charge Coupled Device (CCD) or Complementary Metal-Oxide Semiconductor (CMOS) image

sensor. CCD and CMOS image sensors can be used to provide spectral data from 400 nm up to 1100 nm depending on the sensitivity of the image sensor, typically in wide wavelength bands. Hyperspectral cameras can provide spectral information about an image in bands down to as low as 1.64 nm (Zhang et al. 2012).

Machine vision analysis can make use of a plant's physical and visual features in order to determine its species, based upon a combination of the plant's spatial, spectral, shape and texture properties. More recently, machine vision techniques have included three dimensional data to achieve a more consistent result, based on plant height data alone, or in combination with spectral, shape and texture analysis.

The machine vision literature review presented in this chapter falls into seven broad categories which are:

- segmentation of plant from background;

- identification based on shape;

- identification based on texture;

- identification based on spectral differences;

- identification based on spatial positioning;

- identification based on 3D imagery; and

- real-time implementations.

2.2   Segmentation

Segmentation is a common first step in machine vision analysis for individual plant identification. The purpose of segmentation is to reduce the amount of

data to be analysed by removing all background pixels, therefore leaving only plant pixels. Figure 2.1 displays the spectral properties of a typical green leaf, in which green (approximately 550 nm) reflectance is higher than red (approximately 650 nm) or blue (approximately 450 nm) reflectances. Green plants also exhibit a distinct 'red edge' between red (approximately 650 nm) wavelengths and Near Infra-Red (NIR) wavelengths at approximately 750 nm and higher. This 'red edge' is not as significant in non-green vegetation such as shown in the straws and soil traces in Figure 2.2.
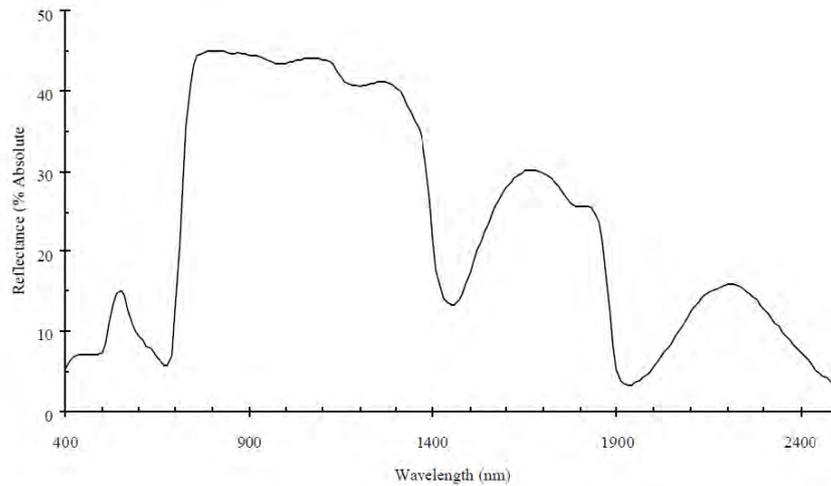
Figure 2.1: Green plant spectral properties from Noble & Brown (2002) displaying the typical leaf reflectance curve between 400 nm and 2500 nm.
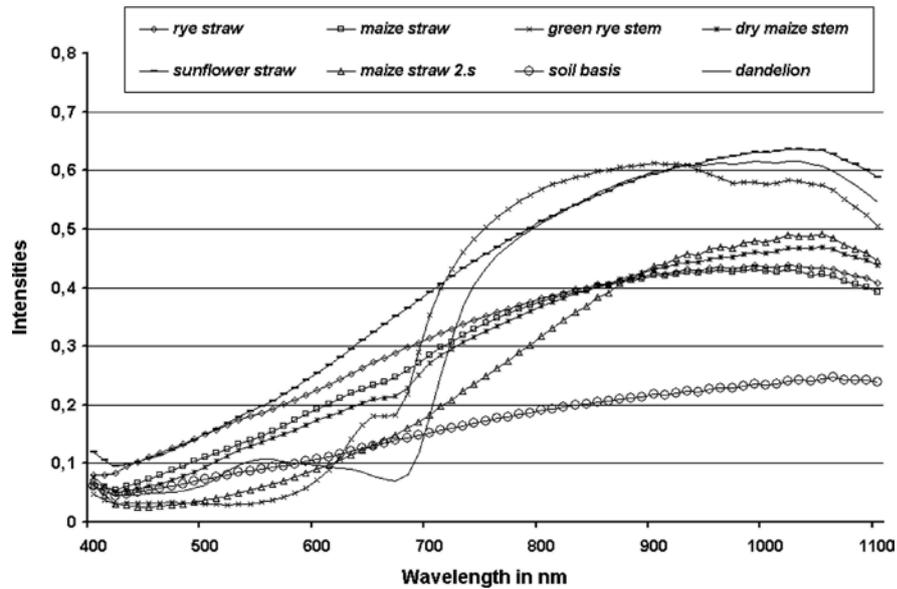
Figure 2.2: Stubble and plant spectral properties from Langer et al. (2006). A distinct red edge is noticeable on the spectrum of green rye and dandelions but not on straw, soil and dry stems.

Piron et al. (2008) found that the best wavelengths for segmentation when determining weeds in carrots were centered on 450 nm, 550 nm and 700 nm which are very close to the wavelengths of an RGB camera at 450 nm, 550 nm, and 650 nm respectively. Typical low cost CCD and CMOS cameras do not have sensitivity above 1100 nm so discrimination of weeds based on reflectance above 1100 nm using typical low-cost cameras is not possible.

Segmentation is a critical step in the image analysis process and the features that need to be isolated in the image determine the segmentation technique used. Segmentation based on colour is a common approach. Woebbecke et al. (1995) evaluated several colour algorithms for segmentation. The evaluated algorithms used the intensity levels of the red (R), green (G) and blue (B) channels of the image and included: $R - G$, $G - B$, $\frac{G-B}{abs(R-G)}$ and $2G - R - B$ (also referred to as 'excess green') and found that the $2G - R - B$ was most effective.

Recent research using the excess green technique includes Swain et al. (2011), who

used excess green to segment images for use in an active shape modeling classification techniques. Jeon et al. (2011) used normalised excess green and an Artificial Neural Network (ANN) for outdoor operation. Sabeenian & Palanisamy (2009) used an offset excess green segmentation with texture analysis and Jafari et al. (2006) applied excess green to segment the image to be used in a discriminant function evaluation. The cited literature shows that excess green (and variants) can offer a computationally fast and efficient means of segmenting plant pixels from background pixels, useful for real-time systems.

Li & Chen (2010) used the Otsu method of thresholding (Otsu 1979) of the hue channel of a HSI image for segmentation. Perez et al. (2000) used the differences between the green and red spectral reflectances by producing a normalised vegetation index and then applying a threshold. Langer et al. (2006) developed a Difference Index with Red Threshold (DIRT) which uses the same spectral bands as the Normalized Difference Vegetation Index (NDVI) but gives a better response on mulched fields.

2.2.1 Illumination effects on segmentation

Reviews on machine vision analysis in real-world settings have found that lighting has been a major factor contributing to errors in classification. Lighting conditions have an effect on the image analysis through changes in intensity, white balance and shadows (Slaughter et al. 2008, Jin & Tang 2009, Noble & Brown 2002, Hong et al. 2012). In outdoor situations the light source can vary from 2000 K in the morning to 8000 K in the shade with white daylight being approximately 5600 K (Lowel Education Centre 2012) dependent upon cloud cover and time of day.

El-Faki et al. (2000) evaluated the effects of illumination and soil moisture as sources of error and found that illumination has a significant effect on correct segmentation, whereas soil moisture did not. Slaughter et al. (2008) found that

illumination was second only to occlusion as a limiting factor for effective machine vision applications. Baron et al. (2002) determined that the 'red edge' could segment effectively but that changing daytime conditions such as cloud cover affected the visible spectrum (VIS) and NIR spectrum differently and this had a significant impact on the accuracy.

2.2.1.1   Alternate colour spaces

In the research literature, there has been significant research effort applied to overcoming the changing nature of outdoor lighting and the degradation on segmentation quality caused to RGB-based segmentation techniques. Applying different colour spaces and algorithms has been evaluated to address changing daytime conditions. Bai et al. (2013) segmented plant matter from background by morphology modeling in the CIE L*a*B* colour space with accuracy of 87% from a dataset of 56 images. Tang et al. (2000) used a genetic algorithm (GA) in the Hue, Saturation and Intensity (HSI) colour space called GAHSI on images with both bright and shadowed regions and achieved similar performance (90%) to cluster analysis-based segmentation on images acquired under uniform lighting conditions.

In summary changing the colour space displayed varying capabilities to segment using alternate colour spaces under differing illumination but all techniques increased the computational complexity compared to excess green or thresholding.

2.2.1.2   Exposure-based effects

Jafari et al. (2006) separated pixels into light and shadow categories in 300 images for classification under outdoor lighting with an average of 85% weed identification. A conclusion from the research was that the lit areas had higher positive classification of weeds than the shadowed areas in the image and that two differ-

ent classifications should be evaluated to improve the result, one for the lit area and one for the shadowed area.

Romeo et al. (2013) developed an expert system which first classified images into two bins; one bin was for images with satisfactory image quality for segmentation; and a second bin was for images of poor quality. The images in bin one had a combined segmentation technique outlined in Guijarro et al. (2011). The second bin was down sampled and had fuzzy clustering applied to it. The results were 91% and 85% respectively over a broad range of varying daylight and the analysis times were on average 0.91 s and 1.61 s respectively, which is not real-time applicable.

Suh et al. (2014) developed a shadow resistant technique which has produced promising preliminary results. The process involves analysing the image to determine if it contains shadows. An excess green segmentation with Otsu thresholding was applied if there was no shadows. If there were shadows, an Otzu's 3-threshold method was applied to the excess green segmentation. The average execution time for the technique was 0.42 s.

Although the exposure based segmentation results are promising the evaluations have not been exhaustive and require significant further testing and validation on real-world no-till images for use in Australian farming systems.

2.2.1.3   Multiple feature extraction and classification

Jeon et al. (2011) developed a system which involved the functions of normalized excess green conversion, statistical threshold value estimation, adaptive image segmentation, median filter, morphological feature calculation and Artificial Neural Network (ANN) in an effort to identify weeds in outdoor illumination. The system had an initial accuracy of 72.6% which was improved to 95% by correcting the error source (incomplete leaves). Limitations in this research were: that all

images were specifically non-occluded, the results for the data were good in full shade, or no shade, however, when the images had both, the outcome was poor in the bright area.

Tian et al. (1997) used an Environmentally Adaptive image Segmentation Algorithm (EASA), to address changing lighting conditions when identifying tomato seedlings at cotyledon growth stage, with up to 78% accuracy.

2.2.1.4   Light-restricting cover

A physical method of addressing variable lighting conditions is the addition of a light-restricting cover such as Figure 2.3.

Figure 2.3: Three metre spot sprayer developed for plant identification and spot spraying at the NCEA (McCarthy et al. 2012). Camera and lights are mounted under white plastic covers with a light-restricting cover around the sprayer.

Wang et al. (2007) found that they needed to retrain their system after daytime changes and that the addition of a light-restricting cover plus lights kept the lighting constant and alleviated the need for continual retraining. The addition of a light-restricting cover to the data collection and spray systems has been

successfully achieved by Hemming & Rath (2001), Lamm et al. (2002), Astrand & Baerveldt (2002) and Tangwongkit et al. (2006).

In reviewing the use of spectral properties in weed detection, Noble & Brown (2002) found that changes in illumination play a major role in segmentation effectiveness, and that controlled, artificial, diffuse broadband light could eliminate a number of the problems as well as permitting night time operation.

2.3   Weed identification based on shape

Morphological features (e.g. leaf shape) are an intuitive choice to be used in machine vision plant identification. However, some species are very similar (e.g. wild sorghum and sugarcane) which may make shape a poor discriminator. Shape matching and classification can be adversely affected by occlusion, varying crop conditions, poor segmentation and image quality (Slaughter et al. 2008).

2.3.1   Effects of occlusion

Franz et al. (1991) attempted to identify partially occluded leaves using curvature to indicate leaf shape. It was found that error was caused by differing leaf serrations and boundaries in multiple leaf images and shape could not be defined.

Tian et al. (1997) reported that occlusion was a significant source of error in applying perimeter, centroid, horizontal and vertical dimensions with an EASA algorithm that targeted tomato cotyledons in different lighting conditions within an individual image.

A standard 'erode' morphological transformation was implemented by Lamm et al. (2002) to address occlusion in identifying cotton plants from grass-like weed with 78.7% accuracy and an acquisition groundspeed of 1.62 km/h.

Gerhards & Chrisensen (2003) used chain code to extract the leaf edge features, then transformed the chain code into a function with a standardised contour and applied a Fourier transform to obtain leaf parameters. The leaf parameters were then compared to a database for classification with an average identification rate of 80% across five plant subgroups identifying occlusion as a major source of error.

Jafari et al. (2006) found that shape analysis was susceptible to occlusions when attempting to combine shape and colour features into a classifier but occlusion of the leaves in the images made the shape features ineffective.

Persson & Astrand (2008) used Active Shape Matching (ASM) to discriminate weeds from sugarbeet, achieving an accuracy of 81-87% highlighting occlusion as a problem for correct classification.

Swain et al. (2011) used an Automated Active Shape Matching (AASM) technique with individual plant segmentation finding occlusion of leaves a source of classification error.

2.3.2   Effects of varying crop conditions

Gliever & Slaughter (2001) evaluated a technique using a radial basis function and a Artificial Neural Network (ANN) and achieved an accuracy of 92%. However, misclassification resulted from hail damage or holes within the leaf and large weed clumps that were similar size to the targeted leaves. Cerutti et al. (2011) applied active contour models to overcome occlusion and found it possible if the model is constrained to fit known leaf shapes.

Sogaard (2005) used ASM on seedlings up to two true leaves with results of between 65-90% accuracy, depending on the weed species, in non-occluded images. The technique Sogaard (2005) used is limited to an operational time window of

several days before the plant outgrows the two-true-leaves growth stage.

Li & Chen (2010) extracted shape features from cotton and weed plants which were then classified with an ant colony organisation algorithm (also used by Dorigo & Blum (2005)) and support vector machine with 94% discrimination accuracy on a limited dataset.

2.3.3   Effects of image quality

Gliever & Slaughter (2001) found that poor image quality from a single CCD camera supplied excessive noise and was not suitable for shape feature extraction.

In a recent review, Copea et al. (2012) found that no one method of leaf classification on its own, is the answer, as plants are diverse in size, shape and colour, and the diversity is both between, and within species. A key conclusion from Copea et al. (2012) was that any technique that relied on a specific set of features to identify plants at a particular growth stage was not robust, as the set of features may not be present at a different growth stage.

2.4   Weed identification based on image texture

Different species reflect light uniquely due to variations in leaf size, shape, angle, reflectivity and clustering. The differences in reflected light create variation in the received colour intensities at the image sensor, producing different textures in the image. Textures have been used for plant identification and individual plant segmentation with differing levels of success. Texture analysis can be categorised into four main groups for feature extraction (Materka & Strzelecki 1998):

1. **Structurally-based**

   Haralick (1979) stated that a structural model of texture is based upon texture primitives (texture elements) and their appearance. Examples of

structural techniques are 'edge per unit area', 'grey level run lengths', 'relative extrema density' and 'relational trees'. These techniques are structural, as their attributes are defined, however, they are also statistical as their spatial connectivity is determined by probability.

2. **Statistically-based**

   Statistical techniques do not define the primitives and their structure. Statistical techniques define the textures' properties by a pixel's probability distribution, and the spatial relationship with other pixel intensities in an image (Materka & Strzelecki 1998). Statistical feature extraction can be separated into two areas being first-order statistics and second-order statistics (Tuceryan & Jain 1998).

   - First-order statistics

     First-order statistics are the probability of occurrence of grey levels in an image at a random location. These statistics can be found with histograms and do not take into account any relationship a pixel's intensity has with its neighboring pixels.

   - Second-order statistics

     Second-order statistics are based upon the relationship between a pixel's intensity and that of the pixels around it, i.e. the properties of pairs of pixels. A popular second-order statistical method is the Grey Level Co-occurrence Matrix (GLCM) (Haralick 1979).

3. **Model-based**

   Tuceryan & Jain (1998) stated that model based methods rely on the construction of a model to define the image. The model parameters contain the qualities of the texture which allow the model to describe the texture as well as synthesise the texture. Types of models used in plant identification have been: fractal-based models (Plotze et al. 2005), stochastic models (Chalak et al. 2011) and random field texture models such as Markov random fields (Sabeenian & Palanisamy 2009).

4. **Transform-based**

   Transform based methods transform the image into a different space whose co-ordinate system represents characteristics other than intensity, such as frequency and size. Types of transforms used for texture feature extraction are Fourier (Ghazali et al. 2008*b*), Gabor (Bossu et al. 2009) and wavelet (Wu et al. 2009) transforms

The GLCM method (a second-order statistical method) is a common tool for feature extraction and has been evaluated by a number of researchers (Meyer et al. (1999), Ghazali et al. (2008*b*), Souza et al. (2008)). Meyer et al. (1999) obtained individual species identification accuracies ranging from 30 to 77% with processing times from 20 to 30 s. It is expected that this processing time would be reduced with modern computer power. Souza et al. (2008) used a limited set of input images with the GLCM and achieved results ranging from 67 to 86%. The literature reviewed indicated that the GLCM is a commonnly used statistical method for plant identification but with variable accuracy.

Burks et al. (2000) evaluated a Colour Co-occurrence Method (CCM) by converting the RGB image to HSI and extracting features from these colour space segments. This approach yielded a classification rate of 90% which was an improvement over Meyer et al. (1999).

Tian et al. (1999) used the Discrete Wavelet Transform (DWT) and the Weed Coverage Ratio (WCR) to detect when there were weeds present in corn and soybean crop. The analysis was based on detecting plant density compared to a threshold (i.e. above the threshold categorised the image as weed present) and were performed in 0.37 s and 0.037 s respectively making the WCR a viable real-time spot spray algorithm, computationally however, the results varied with the background between 47% and 100% which is a high level of inconsistency.

Ahmad et al. (2011) evaluated a Haar Wavelet Transform (HWT) to classify broadleaf from narrow-leaf with an average accuracy of 94% and an analysis time

of 40 ms. Changes in lighting degraded the performance of the HWT and the images contained a single leaf type only. Bossu et al. (2009) evaluated 33 different wavelets against Gabor filtering as a benchmark and found that Daubechies 25 and discrete approximation Meyer wavelets provided better results than Gabor filtering. However, each method required 2.78 s or more to run making it unviable in a real-time spot spray system.

Golzarian & Frick (2011) used Principal Component Analysis (PCA) to differentiate wheat from broome grass and wheat from rye grass with an accuracy of 88% and 85% respectively. The PCA used eigenvalue decomposition and the data was collected from samples grown in controlled conditions in a greenhouse, taken with a high resolution camera (image size of 3648×2736 pixels). Ghazali et al. (2008a) used high and low pass filters to obtain features and developed a feature extraction and classification process called Continuity Measure (CM) which produced a correct result of 98% in discriminating a broadleaf plant from a grass plant in images containing either plant type, but not both.

McCarthy et al. (2012) achieved good results with occlusion tolerance for broadleaf from grass using a line detection technique. A binarised segmentation technique ($G > R$ and $G > B$), followed by a line detection transform and a connected components technique was used to isolate the individual components. McCarthy et al. (2012) used size of component as the discriminator to identify sugarcane with an 85% hit rate and 0.2% false trigger rate.

Local Binary Patterns (LBP) were evaluated by Ahmed et al. (2011) to classify broadleaf and grass weeds with an accuracy of 98% when used with a Suppor vector Machine (SVM) classifier. The evaluation images contained either broadleaf or grass plants, not both, as found in the real-world.

In summary, texture has provided encouraging results for the identification of plant types (i.e. grass leaf and broadleaf) but the computation time for the techniques in the reviewed literature has shown that excessive computation time

is an impediment to most algorithms for use in a real-time spot spray system. The reviewed literature has highlighted limitations in the texture evaluations where the techniques were not evaluated on real-world multi-plant images.

2.5 Weed identification based on spectral differences

Spectral properties (wavelengths) have been used for finding weeds in a fallow situation since 1983 (Haggar et al. 1983). Wang et al. (2001) discriminated wheat from weeds and background, with the most useful wavelengths being 496, 546, 614, 679 and 752 nm. The system performed well where there was a dense plant canopy however performance deteriorated when the plant matter became sparse (<0.02 plant pixels/cm^2). Noble & Crowe (2001) evaluated wavelengths from 250 nm to 2500 nm for weed discrimination and found that 360, 420, 680 and 1930 nm enabled the best discrimination. Noble & Brown (2002) found that illumination changes affects the consistency of the spectral wavelengths.

Borregaard et al. (2000) used two line imaging spectrometers to capture VIS and NIR spectrum with individual identification rates of 70 to 80% on four species at a specific growth stage of less than four-true-leaves. Significant limitations with using spectral wavelengths for weed identification in real-world situations are that the characteristics of the target plant can change, depending upon plant health, weather damage, nutrient and water stress (Slaughter et al. 2008).

Okamoto et al. (2007) used hyperspectral images with Euclidean distance from a validation template as a method of segmentation and achieved a classification accuracy of sugar beet from five weed species of 75% to 80%. Zhang et al. (2012) evaluated the robustness of hyperspectral plant identification in varying seasonal conditions in real-world situations using a line-imaging spectrograph. The wavelengths evaluated were between 400 and 795 nm and the images were

obtained over three years with results of 85%, 90%, 92.7% respectively using a global calibration. Modifying the classifier to calibrate each year individually achieved an identification rate of 95.8% for the year under evaluation.

The reviewed hyperspectral identification systems required relatively slow ground-speeds. Nieuwenhuizen et al. (2010) evaluated spectral properties to discriminate sugar beet from self-sown potato, a speed limitation with the research was the collection speed of 10 mm/s (0.0036 km/h). The system used by Zhang et al. (2012) was operated at a groundspeed of 36 mm/s (i.e. 0.13 km/h). All the literature reviewed, highlights groundspeed as a limitation for using hyperspectral cameras in real-time spot spray systems but newer cameras are improving acquisition times and this will change in the future.

Keranen et al. (2003) evaluated fluorescence as a discrimination technique and reported an image collection time of over 3 s per image which is clearly not suitable for on-the-go field use. Longchamps et al. (2009) found that implementation of a fluorescence system in a real-world situation would have to overcome fluorescence variation due to the growing conditions, plant growth stage as well as plant damage (e.g. hail, pest and disease).

2.6 Weed identification based on spatial positioning

2.6.1 Known row position

The spatial position of the plant in relation to the planted crop row is a method of using prior knowledge to improve classification. If the centre of the crop row can be determined then the assumption that plants centred on the crop row are crop, whereas plants not centred on the crop row are weed can be used and is

discussed below. Weed identification incorporating spatial position can use any of the techniques described in this chapter for segmentation and feature extraction; the classification function then combines prior information about the position of the crop rows with the extracted features.

Wu et al. (2011) used a pixel histogram method to find the centre of the rows of wheat and then highlighted the area in the image from the centre of the crop row out to the edge of the row (on the left and right), therefore plant material not in the highlighted area were weeds. The pixel histogram method had an average weed classification rate of 94% but was only applied to a dataset of five images and the crop rows were unbroken i.e. there were no missing crop plants. Bo et al. (2012) used colour segmentation and row location to identify weeds between the row in outside illumination. The system had an average computation time of 160 ms and in consistent illumination had a 97% weed hit rate. However the accuracy fell to 89% in shadowed areas compared with 92% in strong illumination. This variation in accuracy highlights the issues associated with varying illumination.

Gee et al. (2006) used a double Hough transform to locate rows and region growing segmentation to isolate crop from weed in the rows. Weed infestation rates were then calculated however, occlusion caused the weed infestation rates to be under- or over-estimated which effected overall accuracy. De-Rainville et al. (2014) found the row of crop by applying the Hough transform on binarised images. The boundaries were found and the plants that were found inside the boundary were considered to be a mixture of weed and crop. Features were then extracted and classified to determine weed from crop with accuracies of 90.8%. However, the evaluation data was on small plants (two to four leaf stage) where there is generally less occlusion.

2.6.2   Other spatial techniques

Bossu et al. (2009) used the periodic planting pattern of the crop (both real images of wheat and synthetic images) to identify crop and assumed random weed position. The system is based on Gabor filtering and region-based segmentation applied to a simulated scene.

Berge et al. (2012) developed a patch spraying system called 'Weedcer' which used a calculated Relative Weed Coverage (RWC) to determine weed presence, where the RWC was defined as the weed cover divided by the total plant cover. The young weed leaves and cereal crop leaves were estimated from high resolution RGB images. Weedcer underwent real-time trials with correct classification results of 91% in winter wheat but only travelled at 1.8 km/h with five frames per second analysis. One interesting point was that a Xenon flash was used to overpower the variable daylight lighting.

In summary, the use of prior knowledge based on the spatial position of the plants has been shown to achieve good results and can be a technique that improves classification robustness in real-world situations.

2.7   Three dimensional (3D) imagery

Three dimensional data can be used to provide additional perspectives of a scene to enhance machine vision algorithm accuracies. Monochrome and colour images provide 2D information and the third dimension is depth. Depth disparity maps[1] can be acquired either directly from the acquired data or indirectly by inference from 2D data. Table 2.1 itemises techniques used to obtain depth data found in

---

[1]Disparity map data is the difference in the position of the same point in left and right stereo images. The disparity pixel value is inversely proportional to scene depth (OpenCV Devzone 2013).

the reviewed literature with common techniques being stereo-vision, time-of-flight and structured lighting. 'Active' and 'passive' are two categories referring to the lighting provided for the 3D techniques and Table 2.1 shows that structured light and time-of-flight techniques use active sensing whereas stereo-vision is passive. Each technique is considered further as follows.

Table 2.1: Classification of 3D image techniques (reproduced from Sansoni et al. (2009)).

Classification / Technique	Triangulation	Time delay	Monocular	Passive	Active	Direct	Indirect	Range	Surface Orientation
Laser triangulators	X				X	X		X	
Structured light	X				X	X		X	
Stereo vision	X			X		X		X	
Photogrammetry	X			X		X		X	
Time of Flight		X			X	X		X	
Interferometry		X			X	X		X	
Moiré fringe range contours			X		X		X	X	
Shape from focusing			X	X	X		X	X	
Shape from shadows			X		X		X	X	
Texture gradients			X	X			X		X
Shape from shading			X		X		X		X
Shape from photometry			X		X		X		X

2.7.1 Passive sensing

Passive sensing determines 3D information using ambient illumination of a scene and reflectance from the objects contained in the scene (Sansoni et al. 2009).

2.7.1.1   stereo-vision

stereo-vision is a passive system where depth is found by obtaining positional information from two known viewpoints (two separate images) of the same area. Andersen et al. (2005) used a stereo-vision system with a correspondence technique[2] called simulated annealing to determine geometric properties (height and leaf area) of 10 wheat plants at the five to six leaf stage. Andersen et al. (2005) were able to adequately find the geometric properties with a linear relationship (p<0.001) related to measurements taken of the same ten wheat plants at the five to six leaf stage by a flatbed scanner. However, only ten plants constitutes a limited dataset.

Jin & Tang (2009) used 3D images from a passive stereo-vision system to determine the centre of young corn plants, as 2D images had limitations in particular with overlapped plants (occlusions). The system had a computational time of between 5 and 20 s and was affected by noise from external light. Jeon et al. (2009) developed a technique for passive stereo-vision which improved segmentation under varying illumination; however reduced performance was still evident with high dynamic ranges of illumination.

2.7.1.2   Depth from motion

'Depth from motion' obtains depth information from two consecutive images in a video sequence. If the rate at which an object moves from one frame to the next is known and image capture timing is known, depth can be calculated using correspondence matching as in stereo-vision systems. Sachez & Marchant (2000) used depth from motion and zooming on the camera. Zooming created a greater distance between the same object in both images without having to lower the acquisition frequency (i.e. frame rate) but is not practical in an 'on-the-go'

---

[2]Correspondence is the task of locating the same point in both images. Once this is known, the depth can be found.

system. Sachez & Marchant (2000) also developed algorithms to extract 3D information from a camera in motion to help with crop-from-weeds discrimination. However, unstable features in the images, and low resolution of objects in the image, created significant errors in correspondence matching.

2.7.2 Active sensing

Active sensing is where energy is projected onto the target area and the reflected energy is measured to produce a disparity image. Table 2.1 shows nine active sensing methods however, the research highlighted two: structured light and time-of-flight.

2.7.2.1 Structured light active sensing

Sansoni et al. (2009) state that depth can be derived from the shape of the light pattern reflected off a surface. In a structured light sensing system, a projector or laser, projects light (active illumination) onto the area of interest which is then detected by an image sensor and related to the known illumination pattern. Depth can be determined by interpreting the deformation of the expected pattern in the reflected light, which makes correspondence matching easier than in a passive system (Piron et al. 2008).

Piron et al. (2008) used an active stereo-vision system based on structured light and discriminated weed within a crop of organic carrots, with an accuracy of up to 83%, based on height. Piron et al. (2009) evaluated the benefits of including multispectral data with the height data in the feature set, however only a small improvement of 2% in weed discrimination was found.

2.7.2.2   Time-of-flight sensing

Time-of-flight sensing measures the time taken for a light pulse to travel to the target and back. With an accurate measurement of the time taken, the distance can be calculated. Nakarmia & Tang (2012) used a time-of-flight camera to determine the plant spacing of the crop. To determine the plant spacing, the time-of-flight camera was mounted below the height of the crop in between the crop rows viewing the crop horizontally (i.e. parallel to the ground), as opposed to the standard top viewing configuration of sensors for spot spraying (i.e. above the crop looking down). Nakarmia & Tang (2012) found that the system was not susceptible to colour variations in the crop which was an advantage over 2D identification approaches. The devised system experienced errors when there were multiple plants in the same position, e.g. doubles and triples[3].

Dorrington (2014) stated that current commercial time-of-flight cameras perform poorly in areas where dust is present which may make operation in real-world conditions difficult.

2.7.3   Application of 3D data to segmentation

Paproki et al. (2011) successfully used 3D segmentation in a high throughput plant data acquisition system for plant phenomics research in a laboratory under controlled conditions. The platform combined data from high resolution stereo-vision sensors, multispectral images, infrared images and LIDAR sensors to reproduce a 3D surface mesh overlaid with spectral data with a 9.7% error. However, thirty-two images of the same object on a rotating table were required to perform the analysis and therefore is not readily applicable to real-time spot spraying.

---

[3] A double or triple is where two or three seeds are dispensed simultaneously by the seeder, instead of just one seed.

Seatovic et al. (2008) identified broad-leaved dock (*Rumex obtusifolius*) using 3D images to extract the area of interest and then applied 2D image analysis to the area of interest. A limitation was that the extraction process was up to 10 times too slow for real-time spot spraying. Seatovic et al. (2008) found that the 3D approach improved segmentation performance in areas where 2D approaches failed, such as low contrast images, green-on-green images and noisy images.

Wallenberg et al. (2011) used a fusion of colour and depth information from a Kinect® camera system[4] under controlled conditions to segment individual leaves on a plant using the 'superparamagnetic' clustering algorithm outlined by Blatt et al. (1996). The results showed that a better segmentation result was achieved with the fusion of depth and colour than was achieved by either one individually.

Chene et al. (2012) used the depth and colour images from a Kinect® camera to achieve entire plant phenotyping. The phenotyping was accomplished in a controlled indoor environment with plants that had good height separation, which is not the typical situation in real-world conditions. Figures 2.4 and 2.5 show individual leaves at differing heights being individually segmented even though some leaves are occluded in the colour image. This indicates that the addition of 3D data to 2D data can aid in occlusion tolerance.

---

[4]The Kinect® camera system comprises an RGB camera and structured light depth camera operating simultaneously, used for gesture recognition for the Microsoft X-Box360® and sold individually for robotics.
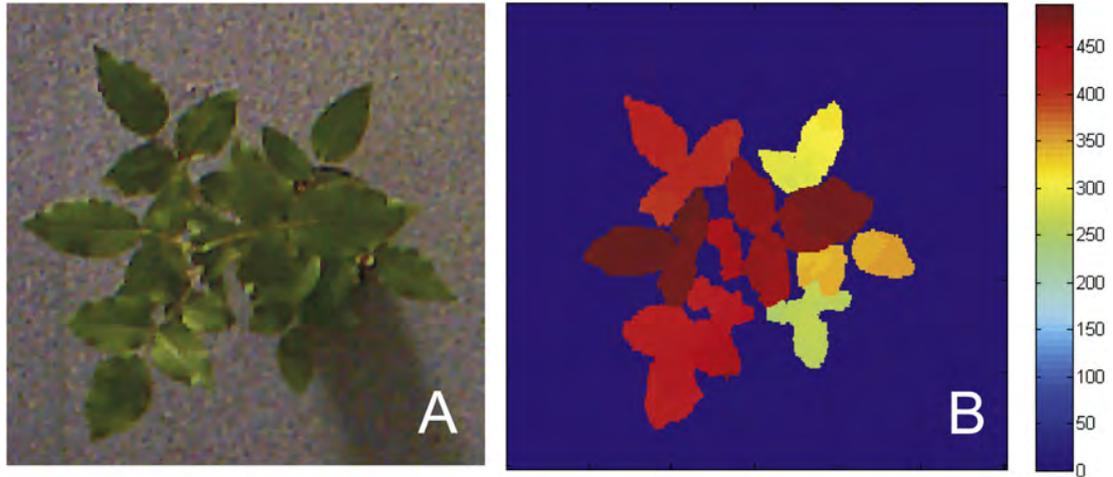
Figure 2.4: (A) Colour image top view of the plant; (B) Depth image of (A) with different heights rendered in different colours, reproduced from Chene et al. (2012).
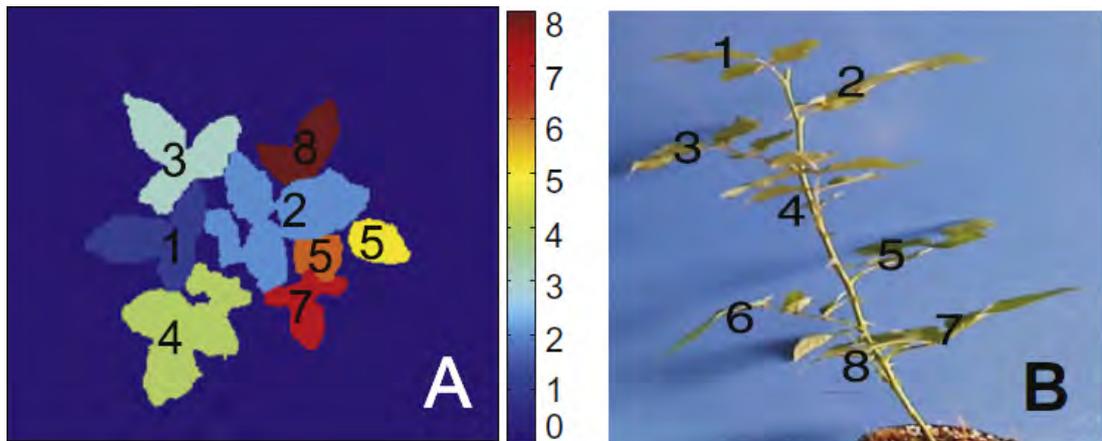
Figure 2.5: (A) Top view of identified leaves with the number of leaves used to identify individual leaf height positions; (B) Side view showing the height position difference between leaves, reproduced from Chene et al. (2012).

2.8   Conclusion – Development of specific research objectives

The foregoing literature review has demonstrated that there are a large number of machine vision techniques which have the capacity to discriminate weed

from crop, but there are multiple limitations. Shape and spectral features are not robust features for classification due to variations in leaf shape from disease, pest and weather damage and changes in growth stage and growing conditions. Texture features, particularly statistical and transform based are commonly used analysis methods for post processed plant identification but are not adequately effective in real-world environment situations in a real-time application. 3D data provided good results for improving occlusion tolerance, however the algorithms were not real-time applicable. 3D data acquisition by stereo-vision cameras proved computationally intensive but structured light cameras were shown to fit real-time applications.

The sources of error that most commonly occur are occlusion and varying illumination (Slaughter et al. 2008, Ji et al. 2009) and successful approaches to overcome these are as follows.

1. Occlusion

   The difficulty in identifying occlusions in 2D imagery is reduced by the addition of 3D image data and the latter has produced promising results Chene et al. (2012). The use of both 2D (colour) and depth images indicated better results than either on their own.

2. Varying illumination

   A number of segmentation techniques have been designed specifically to address the variability of outdoor daytime illumination. However, these segmentation techniques have added complexity and computational time, are not always consistent and have not been demonstrated in a practical implementation in the field. The addition of a light-restricting cover over the camera's field of view, plus a known constant light source, is a successful technique to overcome errors introduced by varying illumination. An additional benefit of a light-restricting cover is the reduction of ambient light that could affect structured light 3D cameras (e.g. the Kinect® camera).

Other significant practical constraints are as follows:

3. Processing time to achieve weed discrimination, overcome occlusions and illumination errors exceeds the time available in a real-time application when operating at commercial groundspeeds.

4. Techniques reported in the literature usually had a narrow operational window (growth stage), which are expected to be impractical for farmers due to varying climatic conditions as well as competing workloads on the farm.

5. The majority of the research was carried out on still photos and not verified for video imagery. Video imagery is required for practical use in real-time.

Therefore the specific objectives of the research reported in this thesis are:

1. To develop an algorithm/s that incorporates 2D (colour) and depth data from video streams to achieve weed discrimination from crop in a real-time, real-world environment at commercially realistic groundspeeds.

2. To demonstrate that the addition of depth data to a suitable image analysis technique can achieve weed discrimination from crop in a commercially acceptable operational window, i.e. at a range of crop growth stages in a real-time, real-world environment.

3. To evaluate the performance of the developed technique under a range of real-world environment conditions; in particular with respect to 3D space versus 2D or depth on their own.

4. To demonstrate that the system is adaptable to a range of crops under practical commercial conditions.

Chapter 3

Field data acquisition

3.1 Introduction

The literature review (Chapter 2) highlighted that the primary sources of error to address in real-world machine vision systems were occlusion and illumination. Accordingly, the data acquisition system for achieving real-time, real-world spot spraying in crops was required to:

- obtain real-time colour and depth data;

- provide consistent illumination; and

- operate at groundspeeds up to 8 km/h in the target crops, namely sugarcane and pyrethrum.

This chapter describes the data acquisition system, data collection plan and agronomic factors associated with the data. It is beyond the scope of this thesis to develop daylight compensation measures for the image sensor. However, the literature review found that an effective method for illumination control is to provide a light-restricting cover and add a constant light source for consistent illumination.

3.2 Data acquisition apparatus

3.2.1 Equipment common to sugarcane and pyrethrum data acquisition systems

Two data acquisition systems were built and deployed into the field for collection of groundspeed, depth and colour data. One system was for pyrethrum crops and the second system was for sugarcane crops. Both data acquisition systems comprised a Kinect® camera system (Microsoft, Redmond Washington USA); a FITPC2® (Compulab, Israel) with a 1.6 GHz dual core processor and 70 gigabyte solid state drive; and 4×2200 lumen 'cool white' LED lights for the pyrethrum unit and 8×2200 lumen 'cool white' LED lights for the sugarcane unit. The lights were powered directly from a 12 volt DC, 70 A hour, deep-cycle battery and the FITPC® was powered from the same battery as the lights, via an SCA 600 Watt 12DC-240VAC inverter.

The Kinect® camera system has a 43° vertical by 57° horizontal field of view with the colour images being 640×480 pixels in size and eight data bits per pixel per channel, from a CMOS image sensor. The depth camera portion of the Kinect® camera system is a proprietary product of Prime Sense Ltd[1] of Israel and has an 11 bit data pixel with 680×480 pixels per image. The Kinect® camera was operated with a fixed white balance and fixed exposure time to ensure consistency of illumination and colour rendition over the collected video footage.

During data collection, the colour image was stored in an uncompressed AVI file, while the depth data was stored in a binary file, and the groundspeed was stored in a text file on the FITPC2®. The data collection on the FITPC2® was achieved by developing a software program using OpenCV (OpenCV Devzone 2013) and Open Kinect (OpenKinect 2013) interface programming libraries to record the images and groundspeed.

---

[1]Prime Sense is a subsidiary of Apple corporation, USA.

3.2.2   Addressing illumination factors

The data acquisition systems were operated either at night, or under a light-restricting cover for daytime use. The light-restricting cover was made from opaque, black canvas, enclosing the viewing area of the camera system so that no direct sunlight could contact the viewing area. The LED lights emitted negligible NIR light (which could have affected the NIR structured lighting system of the depth sensor on the Kinect®).

The brightness of the LED light allowed the exposure time on the Kinect® camera system to be set short to minimise blurring while traveling over the crop. The light intensity at ground level was different for the two data acquisition systems, due to the different mounting heights of the lights and Kinect® camera system, creating differing maximum groundspeeds. Figures 3.2 and 3.9 show the different mounting positions of the camera and lights on each of the data acquisition systems.

Exposure time was determined by trial and error, finding a shutter width integration time register setting of 10 for the pyrethrum system and 50 for the sugarcane system, no unit type supplied. The maximum groundspeed achievable before blur is noticeable in the image at the shutter width settings above, was 10 km/h for the pyrethrum data acquisition system and 5 km/h for the sugarcane data acquisition system. The white balance of the camera was fixed at the values the camera automatically calculated at bootup. Fixing the white balance was important but not the fixed value itself as the white balance could be easily compensated by multiplying a fixed constant value with the red and blue channels.

3.2.3 Groundspeed measurement

The sugarcane data acquisition system prototype used groundspeed from a GPS system. The groundspeed was extracted from the NMEA[2] string output from the GPS. The GPS reciever was a Leica (Leica Geosystems AG, Switzerland) SmartAg with velocity accuracy of 0.03 m/s RMS.

The groundspeed of the pyrethrum data acquisition system was measured automatically by a magnetic wheel pickup sensor with four magnets (Figure 3.1) on the data collection unit. The groundspeed is calculated by the magnetic wheel pickup sensor creating a pulse each time a magnet passes the magnetic wheel pickup sensor when the wheel is rotating; the frequency of the pulses provides the wheel's rotating velocity, in rpm, and groundspeed is calculated by multiplying rotating velocity and wheel circumference together. The accuracy of the magnetic wheel system was evaluated against the GPS system used for sugarcane. The result was that the magnetic wheel pickup varied $\pm$ 0.1 m/s from the GPS.

Figure 3.1: The magnetic wheel speed sensor is shown with the sensor and the magnets mounted on the wheel. Each time the magnet passes the sensor, a pulse is created.

---

[2]The National Marine Electronics Association communication standard for communication with satellite systems (GPS and Glonass). The data is output in ASCII text (http://www.nmea.org/ for more information).

3.3   Sugarcane data acquisition system

The data acquisition system outlined in Section 3.2 was retrofitted to an existing test platform, developed in an NCEA project for the sugarcane industry, that could be towed by a tractor and is shown in Figure 3.2. To be able to use fixed settings for white balance and exposure, the system was operated at night under 8×2200 lumen 'cool white' LED lights, as there was no light-restricting cover fitted to block out the direct sunlight. The sensor (Kinect$^{\circledR}$ camera system) was mounted 1.5 m above the ground to accommodate sugarcane heights up to 1.3 m, and the recorded frame rate was 12 fps at a groundspeed between 2.5 km/h and 5 km/h.
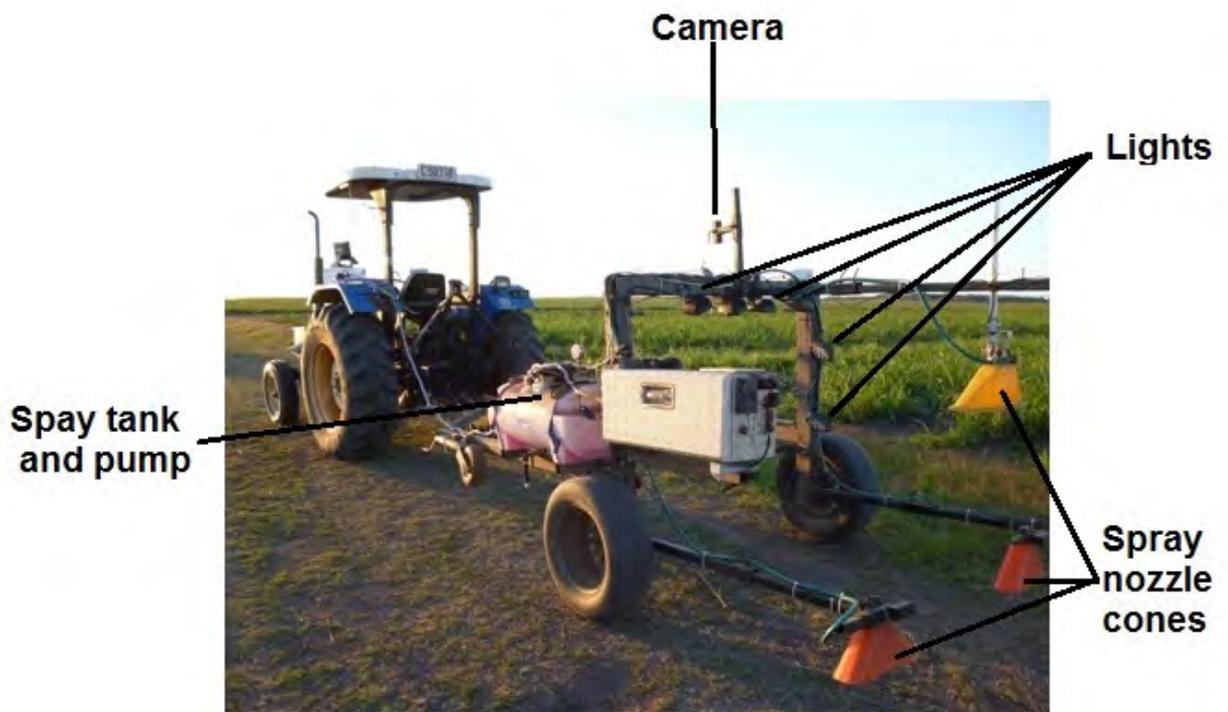
Figure 3.2: Sugarcane data acquisition prototype system.

3.4   Sugarcane data collection

Data was collected from Bundaberg Sugar's farm 'Fairymead' at Bundaberg (latitude -24.791176, longitude 152.356203), Queensland, at night between 6:30 pm and 10:00 pm, between June and December 2012. The data was collected at three growth stages between re-emergence and 1.3 m in height (short 0.1 m - 0.49 m, medium 0.5 m - 1 m and high 1.01 m - 1.3 m). The data was collected from commercial crop varieties grown on 'Fairymead' farm with different ratoons (regrowth crops) and trash conditions. The dominant weeds present in all data collected were guinea grass (*Megathyrsus maximus var maximus* Figure 3.3), nut grass (*Cyperus rotundus L* Figure 3.4) and couch grass (*Cynodon dactylon* Figure 3.5).

Figure 3.3: Typical image of guinea grass (*Megathyrsus maximus var maximus*) cropped from acquired data, hence only low resolution (640×480 pixel images).

Figure 3.4: Typical image of nut grass (*Cyperus rotundus L*) cropped from acquired data, hence only low resolution (640×480 pixel images).

Figure 3.5: Typical image of couch grass (*Cynodon dactylon*) cropped from acquired data, hence only low resolution (640×480 pixel images).

3.4.1   Sugarcane height

Figures 3.6, 3.7 and 3.8are representative of the sugarcane crop at the high, medium and short growth stages respectively. The leaf size can vary significantly between the growth stages and even within a specific growth stage. In approximate terms, the small sugarcane plant category can be up to 0.25 m in height and 0.04 m in width; medium sugarcane plants can be up to 0.8 m high and 0.6 m wide; large sugarcane plants are up to 1.3 m high and 0.1 m wide. Sugarcane can grow significantly higher than 1.3 m in height but for the purposes of traversing a spot spray system over the crop, 1.3 m was the maximum.

Figure 3.6: Sugarcane at a small growth stage (0.25 m).

Figure 3.7: Sugarcane at a medium growth stage (0.8 m).

Figure 3.8: Sugarcane at a high growth stage (1.3 m).

3.4.2   Agronomic factors for sugarcane data

3.4.2.1   The effect of ratoon and trash blanket on results

Sugarcane is a perennial crop and the crop is damaged each year by the effects of weeds, moisture-stress and harvesting machinery. The damage causes the next ratoon of sugarcane to have fewer sugarcane stools per metre. This indicates that the higher the ratoon number, the lower the stool count would be generally, but damage can be variable across the field.

In this research, plant population was estimated for the fields used so that accurate false trigger rates could be determined. Two sugarcane stools per metre (Dart 2013) has been estimated for the sugarcane data from 'Fairymead' farm. This was verified by observation of field conditions during data collection.

The trash blanket in sugarcane is the leftover trash after harvest that is spread

evenly over the ground to provide weed growth suppression. The effects of the trash blanket on the machine vision system is negligible, as the trash does not present itself as standing stubble, there is little difference between trash blanket and non-trash blanket for the machine vision system. In contrast, standing stubble can interfere with the machine vision image capture accuracy by masking the plants from the camera.

3.4.2.2 Optimum sugarcane height for spraying

Data was collected during growth stages of crop that allowed penetration of herbicide to the weeds, i.e. from emergence up to 1.3 m high. Sugarcane can be damaged by the tractor passing over the crop when taller than 1.3 m in height. Professor Bernard Shroeder (Shroeder 2014) asserted that the best spot spray results would be achieved when the sugarcane and guinea grass was at the growth stage between 0.4 m and 0.8 m in height for three reasons:

- Guinea grass can grow from the previous year's root. This provides the regrowth with a significant root system when the plant is small (i.e less than 0.4 m). Therefore, there needs to be sufficient leaf area on which to apply the herbicide such that the plant will uptake a lethal dose of herbicide, sufficient to kill a young guinea grass plant with a large root system.

- The canopy of the sugarcane crop closes over the rows and makes it difficult to deposit the herbicide onto the target weeds outside above the stated growth range.

- A higher groundspeed can be maintained in the stated growth range without damaging the crop.

3.4.3   Sugarcane site criteria

Site selection for data collection was governed by the weed spectrum, soil conditions and trash blanket conditions that was representative of the sugarcane industry in north Queensland.

Site selection

The ideal collection sites were those that covered the main soil types (i.e. black, brown, red and sandy), with each soil type having areas of trash blanket and non-trash blanket.

Collection timing

As stated, data collection was required at three differing growth stages (short, medium and high). The purpose of collecting data from the sites at differing growth stages is to allow effective algorithm development across the growth stages. Sugarcane crops are harvested at staggered times (dependent upon maturity and also harvester availability) from August to December each year. Therefore, it was possible to collect data of varying stages in various fields at the same time.

Data collection speed

Herbicide spray application occurs at groundspeeds from 1 km/h up to 8 km/h for sugarcane. Therefore the data collection system's groundspeed was in this range to provide representative real-world evaluation.

3.4.4   Sugarcane field data

Table 3.1 shows the data data collected from 'Fairymead' farm, Bundaberg, used for development and testing. The collected data comprises data in the identified growth ranges from different fields.

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 com-
puter 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 $\times$ 1.6 m long $\times$ 1 m high) in a crop of pyrethrum March 2012. Author (1.86 m tall) to indicate scale.
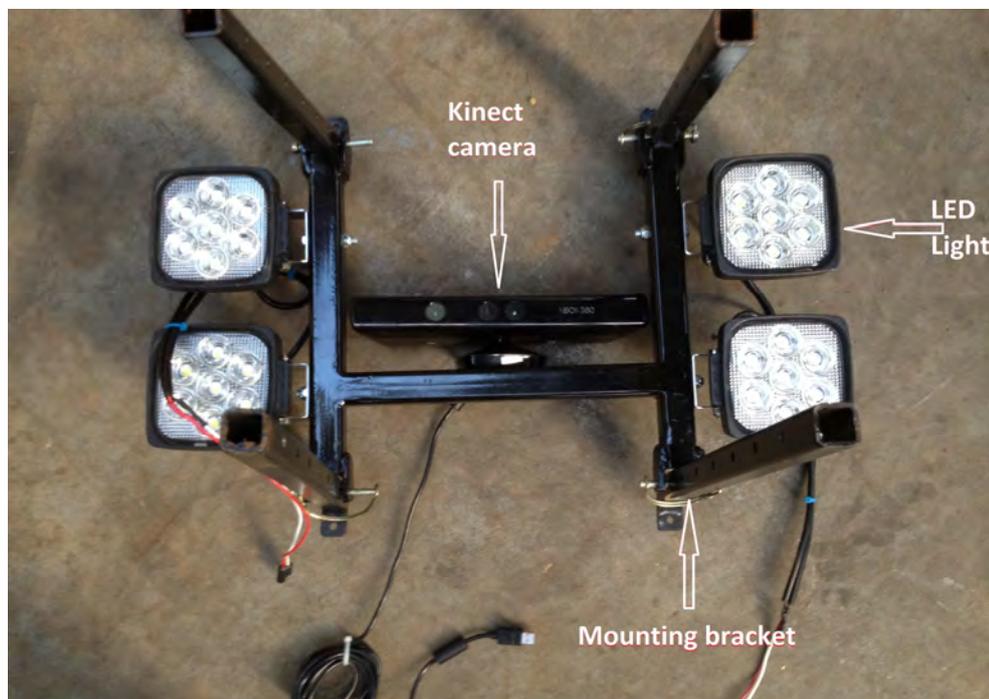
Figure 3.10: Kinect® camera (centre) and lights (four off) for mounting on the pyrethrum data acquisition system. View from below.

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 (m²)**	**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 (m²)**	**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 4

Machine vision methodology, image acquisition and pre-processing

4.1 Machine vision fundamentals

Machine vision can be described as 'the analysis of images to extract data for controlling a process or activity' (Relf 2004) and can be used to automate tasks typically performed by human visual inspection. Machine vision has potential to determine crop from weed in real-time as it allows the analysis of a plant's visual features. This analysis may include a combination of spatial, spectral, shape and texture features, with a large knowledge base existing for machine vision techniques, as shown in the literature review in Chapter 2.

Recently, machine vision techniques have included 3D data to enhance the accuracy and consistency of the result (Section 2.7). Modern machine vision systems may incorporate data for analysis from sources such as visible colour images, near infra red (NIR) images, depth images, multi-spectral images and ultra-violet images.

There are typically five steps in a real-time machine vision system for spot spraying (Figure 4.1) which include three image analysis steps. The five real-time machine vision steps are summarised below and explained in more detail in the remainder of this chapter:

1. **Image acquisition.** The computer acquires the image/s in digital format (i.e. pixel data). Some pre-processing may be required to acquire an image in digital format (e.g. analogue to digital conversion of video signal).

2. **Segmentation.** The analytical process of separating the image into regions that are either 'plant' or 'non-plant'. Ideally, if there is more than one plant in the image, the process will separate each plant for individual analysis.

3. **Feature extraction.** Features are extracted from the plant regions in the image that can be used to distinguish plant species.

4. **Classification.** Classifying the extracted features from the plant regions of the image into unique plant classes (e.g. species or leaf type).

5. **Action and administration.** Action is the control of a physical spot sprayer (spray nozzles). Administration is monitoring and error checking the spot sprayer and physical components of the real-time machine vision system.

Steps 1 and 5 are functions required for a real-time system implementation. Steps 2 to 4 are image analysis steps for plant identification in a real-time machine vision system.
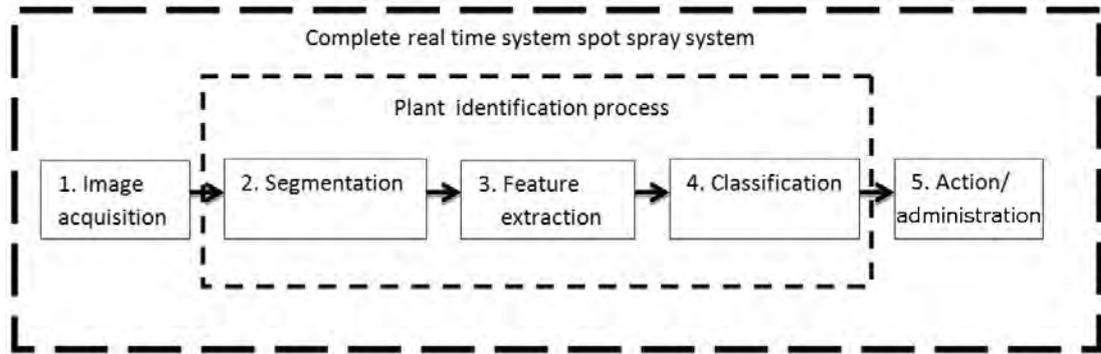
Figure 4.1: Block diagram of the spot spraying system based on real-time machine vision. The plant identification process comprises the inner group.

4.1.1 Machine vision identification architecture

To identify a plant in this thesis, the real-time machine vision system follows the machine vision vegetation identification flow chart shown in Figure 4.2. The flow chart provides the logical flow of decision making, based on image attributes, obtained by the most appropriate machine vision algorithms for segmentation, feature extraction and classification set out in Chapters 5 and 6.

4.1.1.1 Machine vision vegetation identification flow chart

Before following the machine vision identification flow chart decision points, a profile of each targeted species is determined and stored. Each species profile contains image attributes (i.e. features) which relate to characteristics exhibited by that particular species. The attributes are unique to the particular species and can be used to distinguish the species from other plants. The image attributes may be height, colour, shape, texture and size. The species profile provides the criteria to be met at each decision point in the flow chart.

In operation, the real-time machine vision spot spray system sequentially acquires an image and progresses through the processes and decision points set out in the
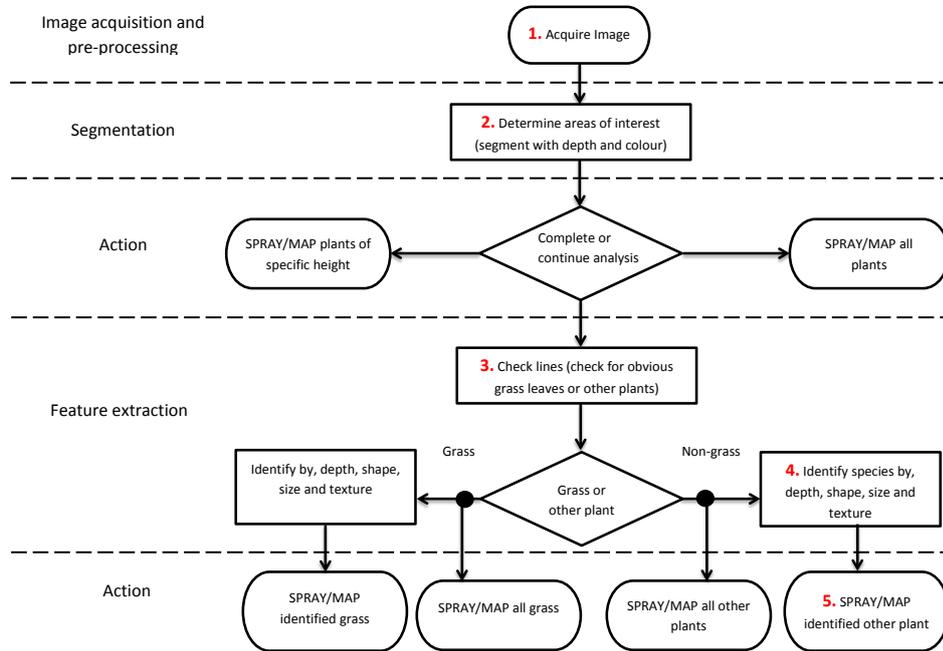
Figure 4.2: Machine vision vegetation identification flow chart setting out the logical flow of processes and decision points based on image attributes. The five numbered points refer to an example of the identification flow of guinea grass from sugarcane.

flow chart of Figure 4.2. As well as processes and decision points in the flow chart, there are output nodes (SPRAY/MAP terminators) in the flow chart where an action can occur and the analysis stops for the current image and waits for the next image. The action can be to spray the weed with herbicide and/or map the plant's position with GPS coordinates for later evaluation as spatial data in an integrated weed management monitoring process.

The flow of analysis through Figure 4.2 to identify a guinea grass plant (as shown in Figure 3.3) is the following. Firstly a colour, size, shape and height attribute profile is created and stored for a target species (e.g. guinea grass) based on a variety of machine vision datasets containing guinea grass as reported in Table 3.1. Guinea grass appears in a clump (non-grass-like) when at small and

medium growth stages. Next, the identification process flows through the stages of segmentation, feature extraction and classification, and more specifically, the list items below which are explained in more detail in Chapters 5 and 6. The number of the action in the following list corresponds to the numbered flowchart items in Figure 4.2:

1. acquire image;

2. determine areas of interest;

3. determine plant category (i.e grass-like[1] or non-grass-like);

4. identify species; and

5. spray or map.

Each of the process points is a part of the machine vision real-time block diagram (Figure 4.1) and highlighted in Figure 4.2. The flow is directed through the non-grass-like decision point because guinea grass initially grows in a clump, and does not exhibit grass-like, leaf-shape attributes in the combined depth and colour features, as does sugarcane. Therefore the identification of guinea grass can be obtained from the non-grass-like features.

Each of the process points are a part of the machine vision real-time block diagram (Figure 4.1). For this example list item 1. is the image acquisition and pre-processing, item 2. is segmentation, item 3. and 4. are feature extraction and classification and item 5. is an action.

4.2   Image acquisition

Video cameras can be operated either in free running (continuously taking and sending images) or one shot/external trigger mode (image is taken and sent when

---

[1]Grasses have narrow leaves relative to leaf height growing from the base of the plant.

a trigger signal is received). The settings for the mode are typically provided in the camera's operation manual. In the case of consumer products (e.g. a webcam), the camera commonly operates in free running mode and the data sent to the computer is handled by the computer's hardware controllers (e.g. the USB controller), controlled by the camera drivers installed with the operating system.

4.2.1 Image acquisition interface

The image acquisition interface drivers used in this research are OpenCV[2] and openKinect[3]. OpenCV is a well known machine vision library and was frequently used in the research literature reviewed in Chapter 2. At the time of undertaking this research, OpenKinect was the only set of drivers that allowed the programmer to customise the driver to access the register settings in the Kinect® camera system which control white balance and shutter speed.

OpenCV is a library of real-time computer vision programming functions developed by Intel®. It is freely distributed, supports C/C++, Java and Python interfaces and runs on the operating systems of Mac, Windows, Android and Linux (OpenCV Devzone 2013).

OpenKinect is a free open source library for interfacing the Windows, Linux and Mac operating systems with the Kinect® camera system (OpenKinect 2013). The Kinect® camera system incorporates a sensor system developed by Prime Sense Ltd.

---

[2]OpenCV `http://opencv.org/`
[3]openKinect `http://openkinect.org/wiki/Main_Page`

4.3   Image pre-processing

Pre-processing, in this research, is concerned with correcting the obtained colour and depth image's geometric co-ordinates, so that analysis techniques that combine the colour and depth data (i.e. the Depth Colour Segmentation Algorithm Section 5.3) have reduced errors arising from the physical camera mounting and timing differences between the cameras. Pre-processing of the depth image, for 3D segmentation, requires the co-ordinate system of the depth image to be remapped to that of the colour image, so that object co-ordinates in the colour image are the same co-ordinates in the depth image. The remapping is performed by following the four steps shown below in Table 4.1. Steps 1 to 3 are outlined in stereo-vision calibration (Bradski & Kaehler 2008) and Step 4 is detailed in Section 4.3.1.

Table 4.1: Steps involved for pre-processing in this thesis.

Step	Operation
1	Distortion correction. The images are corrected for distortions from the lens.
2	Image rectification. Transforms the two images onto a common image plane (The co-ordinate plane in which the image is also reproduced.)
3	Translation. The two images are aligned so that the co-ordinates of each image coincide, perpendicular to the direction of travel. This is necessary because of the horizontal mounting distance difference between the cameras displayed in Figure 4.3.
4	Alignment for groundspeed variation. The two images are aligned so that the co-ordinates of each image coincide, in the direction of travel which is necessary if the two images are captured at different instances in time as the camera is moving.
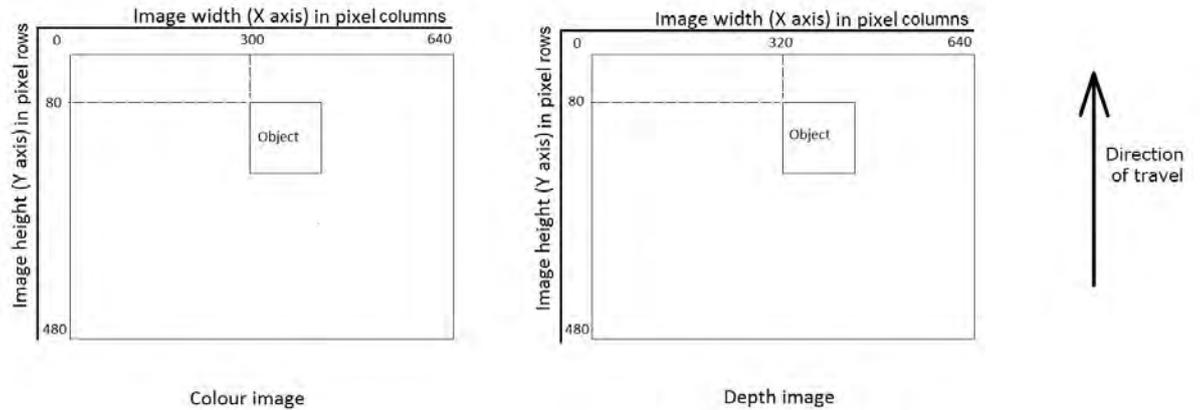
Figure 4.3: Simultaneous image depicting the same object in the colour and depth image co-ordinate systems. The offset in the object X-position between colour and depth images (20 pixel columns in this example) is due to the mounting distance between the colour and depth camera (Figure 4.4).

The depth image represents depth as pixel intensity with pixel intensity decreasing with distance. Once the images (Figure 4.3) are correctly mapped to each other, the depth image intensities are rescaled so that the full scale of depth for the particular mounting height is between 0 and 254 in pixel intensity (zero being the ground surface). Pixels with intensities outside the full scale of depth are referred to as noise and set to 255 in intensity. The images are then made available for segmentation analysis which is the next step in the image analysis process (Figure 4.1).

4.3.1 Steps 1, 2 and 3 – Distortion correction, image rectification and translation

The need to undertake the corrections (items 1 to 3 in Table 4.1) is highlighted in Figure 4.3. Figure 4.3 depicts an object in both the colour and depth image. The object is offset in the object X-position between colour and depth images

(20 pixels) due to the mounting distance between the colour and depth camera. The mounting difference can be seen in Figure 4.4
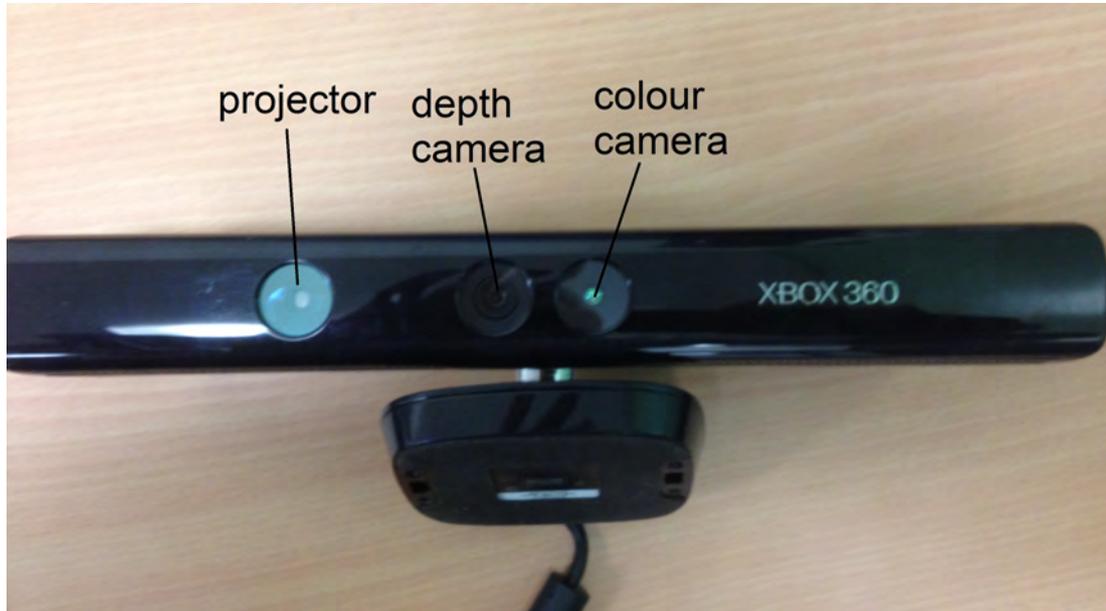
Figure 4.4: Kinect camera identifying the positions of the depth projector, depth camera and colour camera.

4.3.2 Step 4 – Groundspeed alignment

Step 4 (groundspeed alignment) in Table 4.1 corrects the misalignment between the images in the direction of travel. The groundspeed misalignment is due to the time difference between the acquisition of the colour image and the depth image and the distance that an object has moved in co-ordinate location in each image due to the groundspeed of the system. The time difference between the colour and depth image acquisition varies between frames as no two cameras in free-running mode operate at the exact same time intervals due to electronic design and manufacturing differences (e.g. electronic component tolerances).

A visual representation of the misalignment is displayed in Figure 4.5. The object in the colour image starts at row 80 and the same object in the depth image starts at row 100 when traveling at 3 km/h; and row 120 when traveling at 5 km/h.

The object starts in the same column (300) for both the colour and the depth images. The following section outlines the re-alignment function (sometimes called remapping).
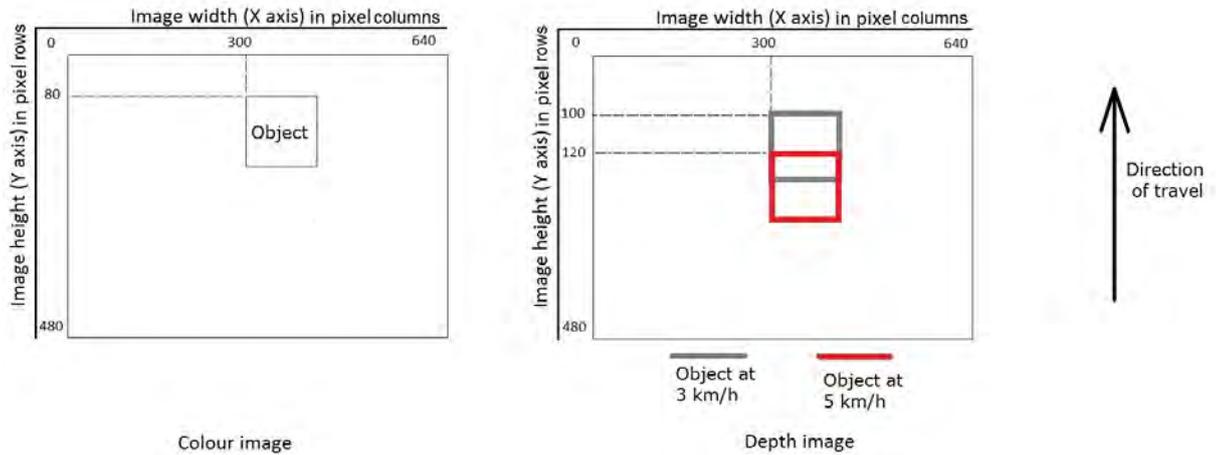
Figure 4.5: Correspond colour and depth images depicting the same object in the colour and depth image co-ordinate systems. The offset in the objects Y-position between colour and depth images (20 pixel rows at 3 km/h and 40 pixel rows at 5 km/h) is due to the time difference between the colour and depth camera image acquisition and groundspeed.

4.3.2.1   Groundspeed re-alignment requirements

Misalignment between the color and depth images due to the groundspeed can be corrected with a re-alignment function. The function achieves re-alignment of the two images by remapping the pixel positions in the depth image so they correspond to the pixels in the colour image. Groundspeed in mm/ms and acquisition time difference in milliseconds between the depth and colour image, is required to accurately re-align the depth image.

4.3.2.2 Groundspeed re-alignment function

To enable real-time operation, a simple alignment algorithm was implemented to realign the depth pixels with respect to the colour pixels based on the following assumptions:

- the motion of the plant itself between the depth and colour image acquisitions is negligible as the time difference is less than 30 ms;

- the difference in plant position in the depth and colour image acquisitions can be corrected by re-aligning the rows of depth image pixels solely in the direction of travel.

Therefore, the columns of pixels in the depth image are remapped to the corresponding co-ordinate positions in the colour image by adding or subtracting an offset based on the distance between the physical mounting of the cameras, and the rows of pixels in the depth image are remapped to the corresponding co-ordinate positions in the colour image by adding or subtracting an offset to the row number (position) in the depth image. The offsets are determined by Formula 4.1 where X is fixed and Y is dynamic.

$$
depthImage_{\text{offset}}(X, Y) = \begin{cases} X = k \\ Y_{\text{offset}} = (t_{\text{difference}} \times s) \div \frac{d}{n} \end{cases} \tag{4.1}
$$

where:

k = A fixed offset correcting for mounting differences.

$t_{\text{difference}}$ = The time difference between the images in ms.

s = The groundspeed in mm per ms.

d = The distance (in the direction of travel) of the image ROI in metres.

n = The number of rows in the image.

4.3.2.3   Factors affecting groundspeed re-alignment accuracy

Leaves that are closer to the Kinect® camera system appear larger than the same size leaf at ground level and can change perspective between the depth and colour images as shown in Figure 4.6. Figure 4.6 shows a high-leaf plant and a low-leaf plant to scale, with a camera at time $t1$ and $t2$ as the camera moves over the plant. The perspective of the plants to the camera can be measured by the angles α and β, the wider the angle, the greater the change in perspective. The image shows that $\alpha > \beta$ highlighting that the perspective on the higher leaf plant has changed more than the low-leaf plant.
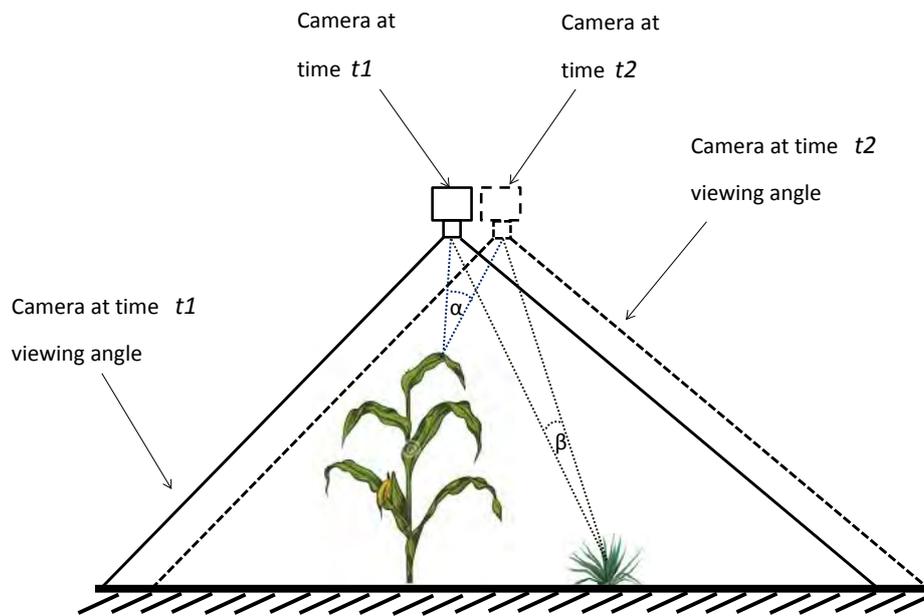
Figure 4.6: Image showing a camera at time $t1$ and $\hat{t2}$. The change in the perspective can be measured by comparing the angles α and β. As shown, $\alpha > \beta$ highlighting a greater change in perspective on the high-leaf plant.

The real-world effect of linear perspective when operating in crops with high leaves (e.g. sugarcane) is that the high leaves in the depth and colour images

don't always precisely overlay each other even after groundspeed re-alignment. An example of the misalignment is shown in Figure 4.7. In Figure 4.7 the high leaves from the depth image are overlayed onto the blue channel of the colour image (the position of the high leaves in the depth image show in the colour image as blue). The superimposed red ellipse labeled 1 shows correctly positioned (overlaid) depth and colour leaves and the superimposed red ellipse labeled 2 highlights incorrectly positioned (overlaid) leaves.

Figure 4.7: Depth image superimposed onto the blue channel in the colour image of sugarcane. The superimposed red circle labeled 1 shows correctly positioned (overlaid) leaves in the depth and colour image. The superimposed red ellipse labeled 2 highlights incorrectly positioned (overlaid) leaves in the images.

A feature of the depth image that helps in minimising the misalignment of high leaves is that objects appear larger in the depth image than in the colour image. The differences between the images is highlighted in colour image Figure 4.8(a) and the greyscale depth image Figure 4.8(b) showing the same leaves in a superimposed red circle. The leaf component highlighted by the superimposed red circle in the images is larger in the depth image, also noise (white pixels) can be seen at the edges of the object in the depth image. Misalignment errors not corrected in this stage are addressed in the segmentation algorithm (Section 5.3).

4.3.2.4   Groundspeed re-alignment options for further assessment

A possible hardware solution available for groundspeed misalignment can be using cameras with a trigger pulse to acquire the image (not free-running). A synchronised acquisition pulse could then be applied to the cameras so that each camera acquires the image at exactly the same time. This method is used in some 3D stereo cameras (e.g. EagleZ® and Bumblebee®). However, this is not the case in consumer depth and colour systems such as the Kinect® camera system.

4.3.2.5   Image pre-processing example

An example of the pre-processing algorithm follows to demonstrate how the algorithm analyses a real-world image. The demonstration compares the sorghum leave's position, relative to a fixed red line in the images. The fixed red line is at the same row number for each image.

1. Figure 4.8(a) is the original colour image with two sorghum leaves ringed in red;

2. Figure 4.8(b) is the original depth image with the same two leaves ringed in red, highlighting a misalignment in the positioning of the sorghum leaves with respect to the red line compared with Figure 4.8(a);

3. Figure 4.9(a) is the re-aligned depth image showing the sorghum leaves at a similar position to Figure 4.8(a) ; and

4. Figure 4.9(b) is the re-aligned depth image superimposed on the red channel of the colour image highlighting the correctly aligned leaves in both images.

(a) Colour image of a sorghum plant in a pyrethrum crop with a red circle highlighting leaves to be compared with the following images.
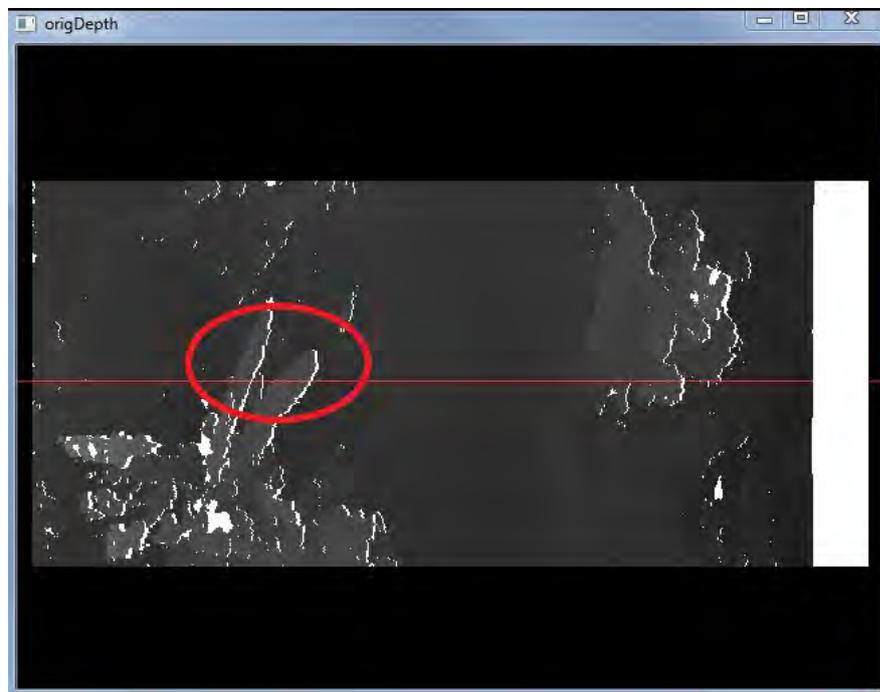
(b) Greyscale depth image of the plant in the colour image above with a red ringed highlighting the same leaves in(a) displaying the misalignment in the Y axis between the two images by comparing the leaf positions to the red line.

Figure 4.8: Images of unaligned colour and depth data.
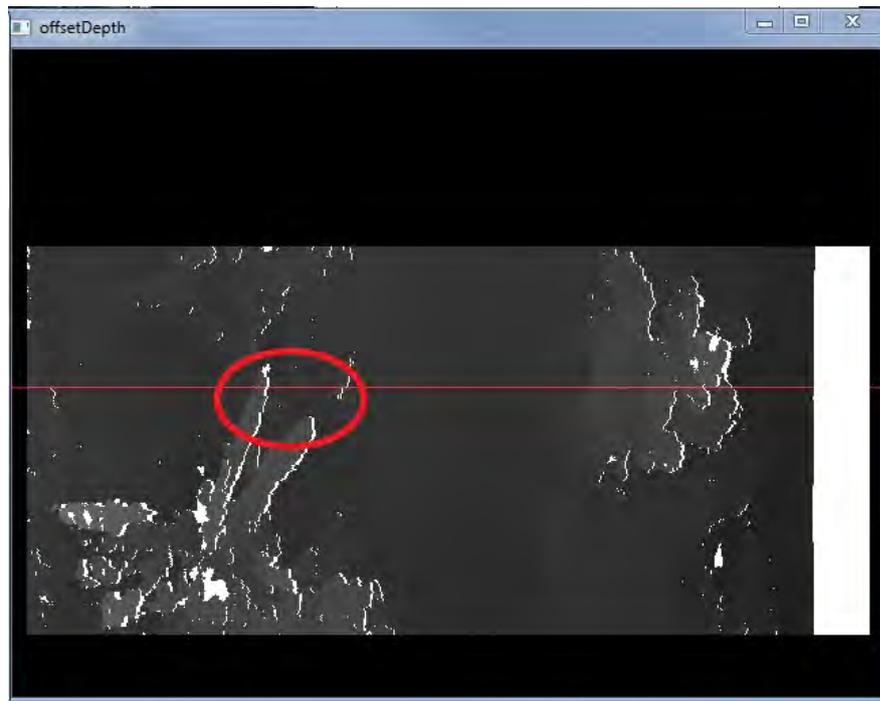
(a) Greyscale depth image of the plant in the colour image with the depth image remapped and the red ellipse highlighting leaves of interest are similarly positioned with respect to the red line.
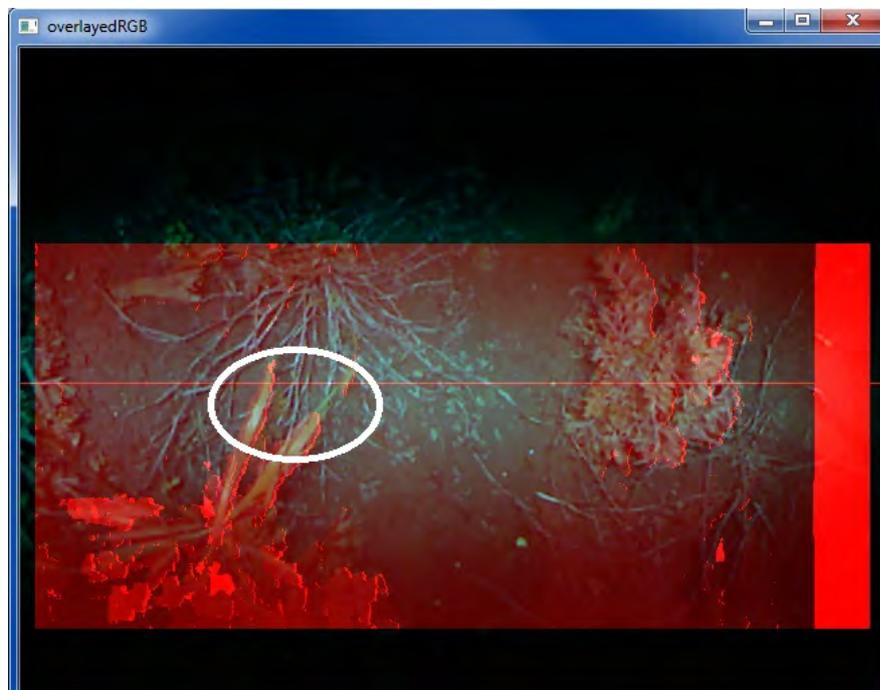
(b) Color image in Figure 4.8(a) with the remapped depth image superimposed in the red channel.

Figure 4.9: Images showing the results of the remapping function. Highlighting the leaves in the colour image are totally covered by the leaves plus noise of the depth image (superimposed on the red channel).

4.4   Summary of Chapter 4

In this Chapter:

1. The machine vision methodology section (Section 4.1) has defined:

   - the machine vision methodology used in this thesis, outlining the five steps of image acquisition (including pre-processing), segmentation, feature extraction, classification and actioning the result, in Figure 4.1; and

   - the weed identification process has been identified and set out in the flow chart Figure 4.2 to provide a decision flow for the machine vision method;

2. The image acquisition section (Section 4.2) has described how the images have been acquired by the use of freeware software tools (openCV and OpenKinect).

3. The pre-processing section (Section 4.3) detailed the added complexity for a machine vision system when two sets of image data from two different cameras are used. This section has demonstrated the effects of added complexities due to mounting distance between the cameras, timing differences between the acquired images and the relationship with groundspeed. Methods have then been described which overcome the added complexities by using openKinect functions for mounting distance between the cameras, and a custom-designed re-alignment function for the timing differences between the acquired images relative to groundspeed.

The corrected images are passed onto the segmentation process (Chapter 5) to separate pixels of value, for feature extraction and classification, from pixels of no value.

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).
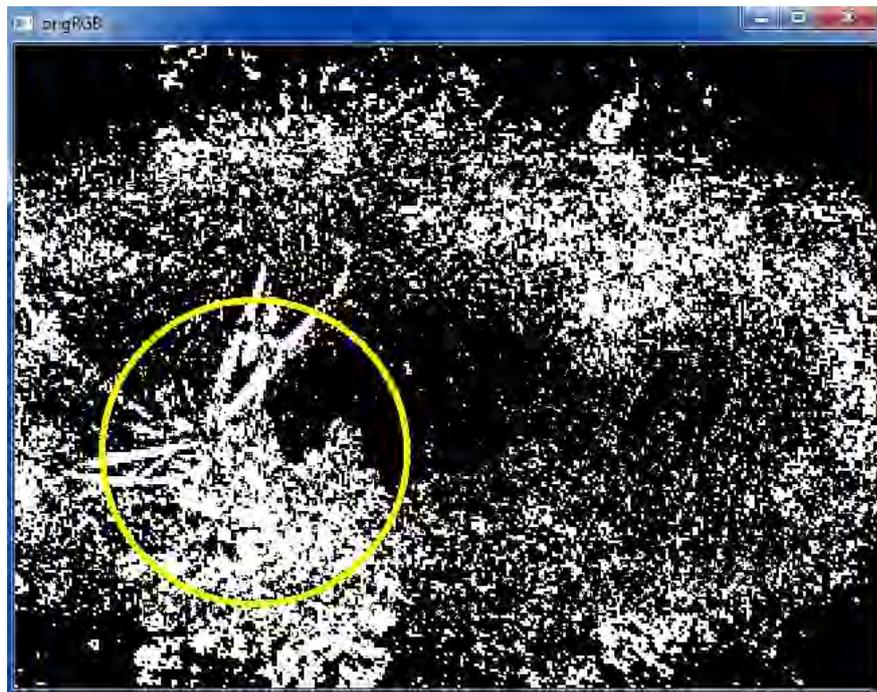
Figure 5.3: Binarised image of Figure 5.1(a) using a BST. Green leaves and dark areas of image are segmented as vegetation.
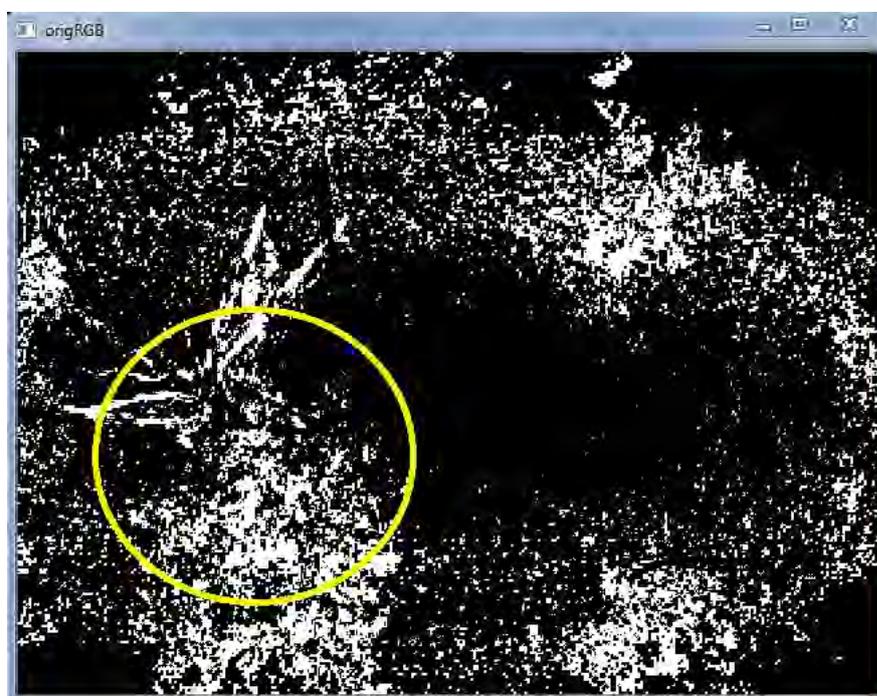
Figure 5.4: Binarised image of Figure 5.1(a) using a BST where R and B are reduced by 10% to lessen the false positives in the image.

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.

	1	
4	seed	2
	3	

Four way connectivity

1	2	3
8	seed	4
7	6	5

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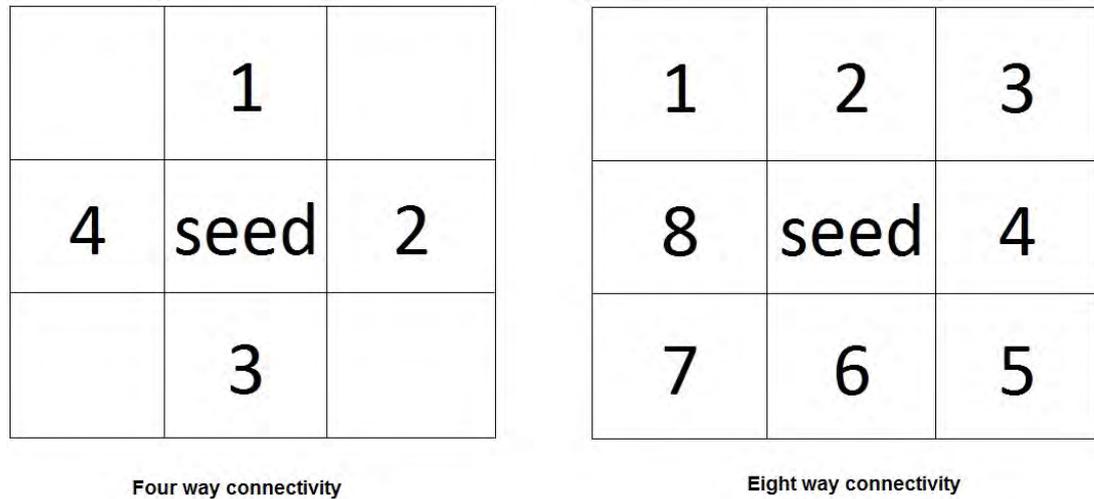
Figure 5.7: Four and eight way connectivity diagrams indicating the directions from the seed pixel that are evaluated.

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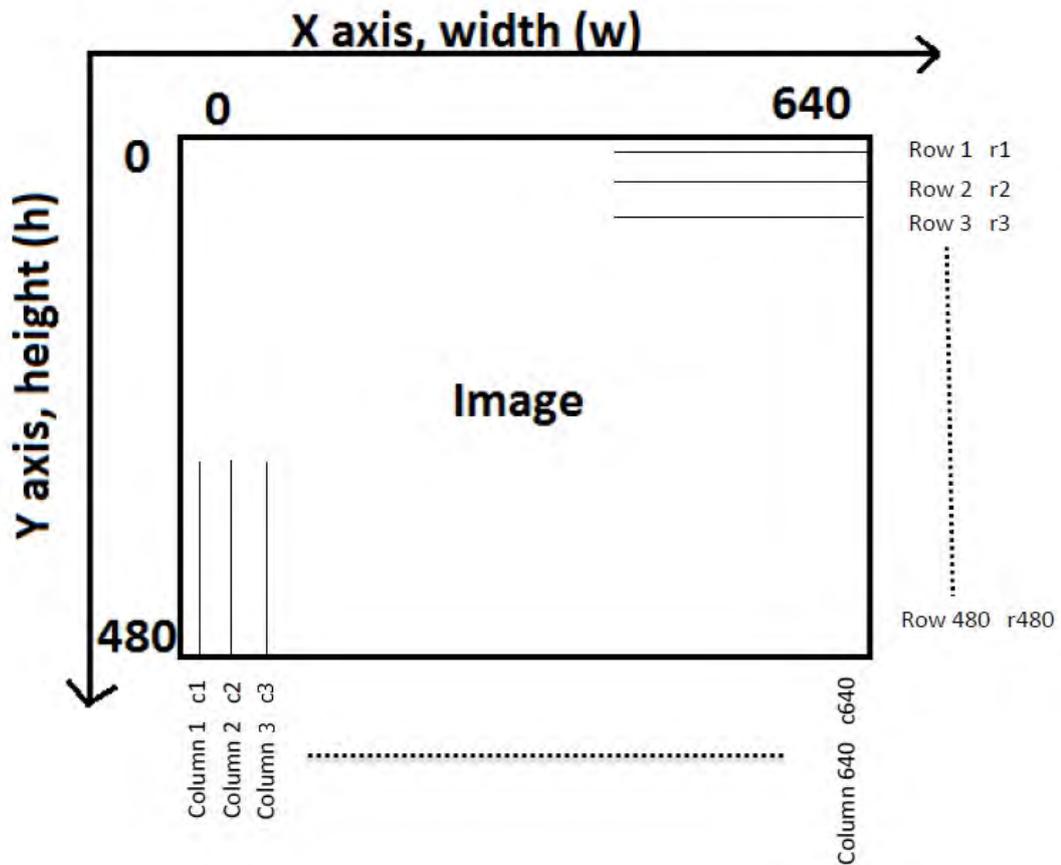
Figure 5.8: Co-ordinate system for the DCSA, with image size 640 × 480 pixels.

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, $connected_{\text{thresh}}$, of the value of the seed pixel.
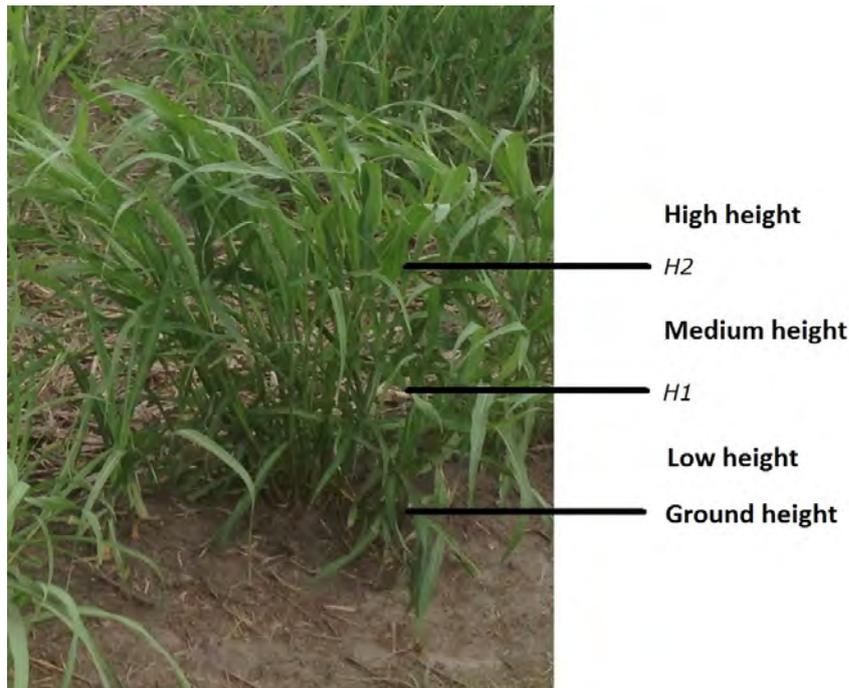
Figure 5.9: Plant height definition for the DCSA classification criteria relative to a sugarcane plant. *H1* and *H2* are user-chosen threshold points for determining low, medium and high areas of a plant.

For example, if $connected_{\text{thresh}}$ is set to 3, then a neighbouring pixel with an value within $\pm$ 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 $\pm$ 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 $percentage\_green_{\mathrm{thresh}}$. The value of $percentage\_green_{\mathrm{thresh}}$ provides the tolerance in the alignment of the depth and colour images, i.e if $percentage\_green_{\mathrm{thresh}}$ 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 $variation_{\mathrm{max}}$ imposes a limit on the absolute variation between a neighbouring pixel's intensity and the

first (original) seed pixel value of $\pm\ variation_{\max}$.

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). Eval-
uation 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 atten-
tion 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 Bina-
rised 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 regrowth) 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).

Guinea grass was observed to grow in clumps compared to sugarcane which has alternate leaves around a main stem. The leaves of the guinea grass were a different size and shorter than the sugarcane leaves. The collected field data showed that at differing times the guinea grass exhibited a yellow emphasis of the green portion of the spectrum (e.g. centred on typically 580 nm) where the sugarcane was more centred in the green spectrum (e.g. typically 540 nm). It is recognised that the variation in colour would be inconsistent depended upon variety, growth stage, nutrient deficiencies and moisture excess or deficiency (Shroeder 2014).

5.6.2   DCSA setup and operation in sugarcane

5.6.2.1   DCSA setup parameters for sugarcane operation

The components for the retained image after application of the DCSA are those components:

- whose maximum height fall below *H2*;

- are green;

- are above a threshold size; and

- are not long and thin.

The parameters required to achieve the retained image results are set out in Table 5.3.

Table 5.3: DCSA setting identified in Section 5.3.3 for sugarcane. X signifies 'don't care'.

Setting / flag	Value
DEPTHSTATS	true
COLOURSTATS	false
$total\_size_{\mathrm{thresh}}$	300
GREENCHECK	true
ALIGNMENT_ERROR	X
$percentage\_green_{\mathrm{thresh}}$	X
REDUCE_LINES	true
$RLMin_{\mathrm{thresh}}$	0.4
$RLMax_{\mathrm{thresh}}$	1.6
$RLSize_{\mathrm{thresh}}$	0.3
KEEP_LOW	true
KEEP_MEDIUM	false
KEEP_HIGH	false
$H1$	40
$H2$	60
MAXLENGTH	true
$variation_{\mathrm{max}}$	100
$connected_{\mathrm{thresh}}$	3

5.6.2.2   DCSA operation in sugarcane

The DCSA sorted the higher leaves (associated with sugarcane) into the deleted image as defined in Section 5.3.3.5 (KEEP_LOW flag set to true). The remaining components that are not green (GREENCHECK flag set to true), or below the threshold size ($total\_size_{\text{thresh}}$ set to 300), or that are long and thin (RE-DUCE_LINES flag set to true), are also sorted to the deleted image. All other components are sorted into the retained image. The retained image was then overlaid with the RGB image data for the identical pixel locations and used for feature extraction and classification, which may be more computationally intensive, but is on a smaller data set.

Figures 5.12 and 5.13 are a visual representation of the segmentation process for a typical sugarcane/guinea grass occlusion. Figure 5.12(b) shows the BST which highlights the issue of undetected occlusion in the resultant image Figure 5.12(c). In contrast Figure 5.13 shows the DCSA output which demonstrates the occlusion tolerance capability of the DCSA.

(a) Original colour image of sugarcane and guinea grass.

(b) BST of image (a).

(c) Connected components applied to the binarised segmentation technique of image (b) with incorrect segmentation of sugarcane and guinea grass.

Figure 5.12: Image sequence showing original colour, binarised and traditional connected components of binarised image. Red ellipse highlights sugarcane plant and guinea grass. (Analysis continues in Figure 5.13 following).

(a) The DCSA technique applied to Figure 5.12(a) showing segmentation of sugarcane and guinea grass.
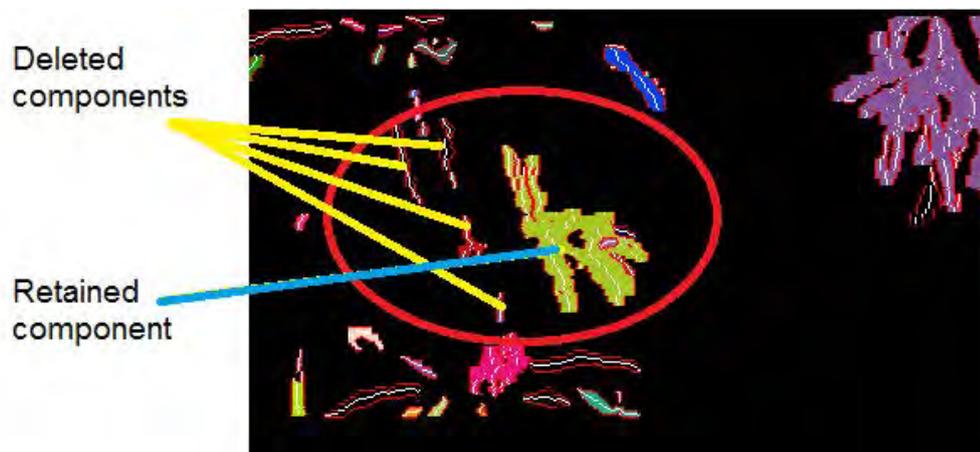
(b) Retained components in block colour and deleted components represented by white lines.

(c) Original image (Figure 5.12(a)) overlaid with retained components.

Figure 5.13: Image sequence showing results for the DCSA function applied to the original image of Figure5.12(a). Red ellipse highlights a sugarcane plant and an occluded guinea grass plant.

5.6.3 Results and discussion for the DCSA used in sugarcane

5.6.3.1 Evaluation data and terminology

Results for the evaluation of the DCSA in sugarcane have been obtained from 856 pairs of frames (colour and depth) from sequential video data collected on 10/10/2012, field 13-A (Table 3.1). This data was considered typical, containing guinea grass and sugarcane within the optimum range of growth stage for spraying (Section 3.4.2.2). The soil was black to grey with no significant trash. The data contained 22 guinea grass plants; and approximately 40% of the frames also contained nut grass ranging from sparse (i.e. less than 10% of ground area in image) to complete coverage of the ground area in the image. The size of the weeds in the data used in Table 5.4 is defined as follows and illustrated in Figure 5.14:

- Small - A guinea grass plant fitting in a circle of between 0.1 to 0.2 m diameter and occupying at least 30% of the circle, by pixel area.

- Medium- A guinea grass plant fitting in a circle of between 0.2 to 0.4 m (inclusive) diameter and occupying at least 40% of the circle, by pixel area.

- Large - A guinea grass plant fitting in a circle of between 0.4 to 0.8 m diameter and occupying at least 40% of the circle, by pixel area.

Figure 5.14: Image of medium weed inside a 0.3 m circle approximately 40% filled.

5.6.3.2   Results for occlusion in sugarcane

The DCSA analysis function was applied to the 856 frames of test data and the resultant retained image and evaluation data (Section 5.5.2) for the retained image after classification were recorded. In addition to the manual assessment for occlusion, each frame was manually assessed for the size of the guinea grass being occluded. Table 5.4 displays the results of this analysis.

Table 5.4: Occlusion tolerance of the DCSA in sugarcane.

Guinea grass growth stage	Number of plants in video data*	Number of plants isolated* in the video data	Number of plants occluded* in video data	Number of occluded plants after BST (%)	Number of occluded plants after DCSA (%)
Small	10	3	7	7 (100%)	0 (0%)
Medium	11	4	7	7 (100%)	0 (0%)
Large	1	1	-	0 (0%)	0 (0%)
Total	22	8	14	14 (100%)	0 (0%)

* obtained by visual inspection

Table 5.4 shows that 14 of the 22 guinea grass plants were occluded in the dataset. The DSCA successfully segmented all of the occluded guinea grass plants in small and medium growth stage categories. The BST results for images with occlusion were poor as the technique could not segment any occluded guinea grass plants. There were 320 frames containing guinea grass and only 1 frame where the DCSA failed to segment the guinea grass from the sugar cane. Hence the failure to segment rate for guinea grass is 0.31% or an accuracy when occluded of greater than 99%.

5.6.3.3 Results for sorting in sugarcane

Reduction in the data requiring feature extraction and analysis can be determined in two ways, namely component reduction and total pixel number reduction. Table 5.5 shows that the reduction in components falls from an average of 28 components per frame down to 12 components per frame, i.e. a 58% reduction per frame. The standard deviation highlights a variation between frames for compoents and pixels which would occur depending upon the number of plant pixels in the image.

Table 5.5: Reduction of components of the DCSA in sugarcane; statistics of experimental results.

Statistic	Average	std deviation
Number of depth components per image	28	13
Number of depth components after segmentation	12	6
Percentage reduction in the number of depth components	58	-
Computation time of DCSA in ms	20.1	0.8

Table 5.6 highlights the reduction in pixels of 84% per frame after DCSA segmentation.

Table 5.6: Reduction of pixels from the DCSA in sugarcane.

Statistic	Average	std deviation
Depth component pixels per frame before segmentation	27096	12006
Depth component pixels per frame after segmentation	4297	3269
Percentage reduction in depth component pixels	84%	-

Table 5.7 shows a reduction of 76% in the number of pixels for the DCSA after sorting compared to the number of pixels in a BST. Again, the standard deviation highlights a variation between frames for components and pixels which would occur depending upon the number of plant pixels in the image. The difference in the number of total pixels between the BST (17774) and DCSA (27096) before sorting is due to the effect of depth components appearing larger than the colour components (Section 4.3.2.3).

Table 5.7: Reduction of pixels from the binarised segmentation technique in sugarcane.

Statistic	Average	std deviation
Number of pixels per frame after BST segmentation	17774	9396
Reduction retained depth pixels compared BST pixels	76%	-
Computation time of BST in ms	0.74	0.18

5.6.3.4   Real-time application in sugarcane

The average execution time of the DCSA for the 856 image pairs was 20 ms with a standard deviation of 0.8 ms (Table 5.5). The execution time required for weed spot spraying at commercial groundspeeds is discussed in Chapter 7 and is 23 ms (average time plus 3 standard deviations, i.e. 99% of data and would exclude outliers). The DCSA execution time for the segmentation technique satisfies this requirement. Additionally, the ability of the DCSA to sort components

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
$variation_{\max}$	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.

Chapter 6

Feature extraction and classification

6.1 Introduction

Feature extraction is the third step in the real-time machine vision spot spray system (Figure 4.1) and is applied to the segmented plant regions (components) in an image. Features chosen for extraction in this thesis were those that produced greatest separable results for the categories of crop and non-crop. The more separable the extracted features are, the simpler and more robust the classification of the segmented plant regions into the different plant categories can be. Feature extraction methods can be a considerable computational load on the CPU depending on the mathematical complexity of the feature. Therefore, the computational requirement of the feature extraction method must be evaluated in real-time implementations.

The fourth step in the real-time machine vision spot spray system after features are extracted is to classify components into different plant categories based on the extracted features. This chapter evaluates machine learning techniques that

have provided satisfactory results (as reviewed in Chapter 2). These are Support
Vector Machine (SVM) (Ahmed et al. 2011), neural networks (ANN) (Jeon et al.
2011), k-Nearest Neighbour (k-NN) (Astrand & Baerveldt 2002) and Naive Bayes
(Zhang 2004).

Machine learning techniques can be supervised and unsupervised. An unsuper-
vised technique is where the data is grouped into similar data groups but the
data groups are not related to any particular plant category (e.g. region growing
segmentation). A supervised technique is where a training set of data representa-
tive of the categories for classification is used to train the classifier and then test
data sets are used with the classifier to label the data into the trained categories
(Bradski & Kaehler 2008).

This chapter documents the development process and evaluation of feature ex-
traction and classification techniques for sugarcane and pyrethrum. Common
techniques are implemented and evaluated, with resultant low accuracy, there-
fore new techniques are developed, implemented and evaluated.

6.2 Overview of existing feature extraction and classification techniques

Visual inspection of images from sugarcane and pyrethrum crops indicated a tex-
tural difference between weeds and crop. To determine if texture would be a
satisfactory means of discrimination between weeds and crop, an evaluation pro-
cess was carried out on texture analysis methods that gave satisfactory results
according to the literature review (Chapter 2). The evaluated methods were Grey
Level Co-occurrence Matrix (GLCM), Grey Level Run Length Matrix (GLRLM)
and Local Binary Patterns (LBP). A brief outline of the different texture extrac-
tion techniques follow in Sections 6.2.1 to 6.2.3; Section 6.2.4 sets out a method
for evaluating the separability of the texture features and Section 6.2.5 to 6.2.8

presents classifier techniques.

To be effective in a real-world situation, feature extraction is required to be both rotationally invariant and illumination invariant, because the presentation and orientation of the plant to the camera varies from frame to frame as the camera moves over the plant. Illumination variance in this thesis is minimised either by enclosing the region of interest in a light-restricting cover (hood) with a known constant light source, or by operating the system at night with a known light source. Rotational invariance is provided via the software implementation, and is described further within each technique below in Sections 6.2.1 to 6.2.3.

6.2.1   Grey Level Co-occurrence Matrix (GLCM)

The GLCM technique outlined in Haralick (1979) is a statistical texture feature extraction technique. Hall-Beyer (2013) defines the GLCM as "a tabulation of how often different combinations of pixel brightness values (grey levels) occur in an image". The GLCM addresses rotational variance by assessing and averaging the combinations of pixel brightness in four different directions ($0°, 45°, 90°$ and $135°$). The OpenCV implementation of the GLCM was evaluated and provided the following features: contrast, homogeneity, entropy, energy, correlation, cluster tendency, cluster shade and maximum probability.

6.2.2   Grey Level Run Length Matrix (GLRLM)

The GLRLM technique, also outlined in Haralick (1979), determines a texture's statistical feature based upon the number of contiguous pixels with similar or the same grey levels in a particular direction. A coarse texture would have a large number of contiguous similar grey levels and a fine texture would have only a few (Haralick 1979). A C/C++ implementation was adapted from Matlab$^{®}$ (Math-Works Inc, Ismaning Germany) which provided the features described in Tang

(1998). The features were: short run emphasis, long run emphasis, run length non-uniformity, run percentage, grey level non-uniformity, low grey-level run emphasis and high grey-level run emphasis. Rotational invariance was obtained by calculating and averaging statistics in the principal axes (0° and 90°).

6.2.3 Local Binary Pattern (LBP)

Ojala et al. (2002) stated that LBPs are found in greyscale images and represent the spatial structure of local areas in an image. The pixel patterns fall into two categories, 'uniform' and 'non-uniform'. Uniform pixels are defined as pixels that contain useful information for describing a texture (e.g. edges), and non-uniform, supplying no useful information (Ojala et al. 2002). LBP uniform and non-uniform pixel description and implementation of rotational invariance can be found in Ojala et al. (2002). The notation for the LBP used in this thesis is $LBP_{8,1}^{riu2}$ which indicates rotationally invariant, less than two transitions, eight neighbours, with a radius of 1.

The LBP was chosen as the functionality was able to be efficiently implemented in a real-time system. The real-time implementation of the LBP in this thesis is shown in Appendix F. In contrast to the implementation of the LBP outlined in Ojala et al. (2002), the real-time implementation used in this thesis is novel in its approach of comparing and counting pixels, and pixel transitions as opposed to computationally time consuming multiplications. It is evaluated in this chapter.

6.2.4 Receiver Operating Characteristic (ROC) Curves

ROC curves are a feature selection tool to aid in identifying the features that will be most successful in a classifier. ROC curves are a description of diagnostic accuracy achieved by generating a plot of the false-positive rate versus the true-positive rate while an input variable is altered (Metz 2006). The result of each

feature is given as a value from 0 to 0.5 with the feature closest to 0.5 providing the greatest separability of data and hence the greatest opportunity for classification. Other tools are available (e.g. a Student 'T' test) however, the output of the T test is either acceptable or unacceptable with respect to a specific null hypothesis but doesn't provide the 'degree' of how good a feature is. The ROC curve is a tool that does provide the degree of suitability of a feature with a working implementation in Matlab.

6.2.5   Support Vector Machine (SVM)

A SVM is a supervised machine learning classifier, defining any two points in a higher dimensional space to make them more likely to be linearly separable. In the learning phase, the algorithm learns the boundaries of the data categories in higher dimensions that create the maximum category separation (Bradski & Kaehler 2008).

6.2.6   Multi Layer Perceptron (MLP)

The MLP is a neural network discriminative algorithm that typically has hidden units between input and output nodes to better represent the input signal (Bradski & Kaehler 2008). The hidden units are a network of non-linear elements with input to the non-linear elements being a weighted sum of outputs of previous elements, and are typically thresholded (Nilsson 1998). Neural networks can be slow to train but are computationally fast to execute (LeCun et al. 1998).

6.2.7 K-Nearest Neighbor (k-NN)

The k-NN is one of the simplest discriminative classifiers. Training data is stored with labels. Thereafter, a test data point is classified according to the 'majority

vote' of its K (distance) nearest other data points (in the Euclidean sense of nearness). The k-NN is often effective but is computationally slow in execution and can require a considerable amount of memory (Grossberg 1987).

6.2.8   Naive Bayes

The Naive Bayes classifier is a probabilistic graphical model (Zhang 2004) and is the application of Bayes theorem with naive assumptions. The naive descriptor refers to the input assumption that all features are independent. Bayes theorem details the probability of an event, from data that may be related to the event.

6.3   Evaluation of existing feature extraction and classification techniques

Images were manually segmented into smaller sub-images of weed and crop for preliminary testing of the feature extraction algorithms outlined above. The sub-images were sampled by means of a sliding window[1] outlined in Section 6.3.1. The training and test set both consisted of sub-images of weed and crop in proportions set out in Table 6.1. For sugarcane, the crop images were sugarcane and the weed images were guinea grass. In pyrethrum, the crop images were pyrethrum and the weed images were various species common to all collection sites (Table 3.8).

The outputs of the feature extraction algorithms were stored in a Comma Separated Values (CSV) file, and further evaluated to identify the most effective features for discrimination. A ROC curve was used to identify the most effective features. The features identified by the ROC curves were stored and applied to classifiers for identification. The evaluated classifiers were the SVM, MLP, k-NN

---

[1]A sliding window sub-samples the data into small regions for evaluation of that region with the window moving across and down the image.

and Naive Bayes.

6.3.1 Software development for evaluation of existing techniques

The evaluation software for the existing feature extraction and classification techniques was written in C/C++ using the OpenCV library (Section 4.2.1). The same library that was used in the real-time implementation of the techniques. Therefore, the outcomes of the evaluations would be more likely to be replicated, in the real-time implementation, as differences caused by alternative implementations of the feature extraction and classifier techniques using various other languages, development environments, compilers and or libraries would not occur.

Matlab® was used for identifying the feature with the highest separability, because there was a Matlab® implementation of ROC curves available, whereas OpenCV did not provide such an implementation available. A block diagram of the three software applications written for the evaluation of existing techniques is shown in Figure 6.1.
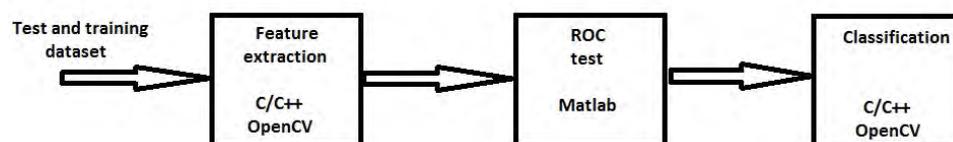
Figure 6.1: Block diagram of software written to evaluate the existing feature extraction and classification techniques.

The first block applied the texture analysis methods to a test image and stored the results in a CSV file. The texture analysis process was iteratively applied to the image by means of a sliding window, with the size ranging between 3 and 29,

incrementing by two at each iteration. This block was written using C/C++ and OpenCV.

The results of the texture analysis techniques were loaded into a second block, written in Matlab®, which applied the ROC test to the texture analysis results. The three most effective features from each feature extraction technique, their respective ROC value, and the window size, were then stored in a CSV file.

The results from the Matlab® block were then read into the classification block which applied the three most effective features from each feature extraction technique separately to each of the classification techniques and stored the results in a CSV file. This block was written using C/C++ and OpenCV.

6.3.2 Results of existing classification techniques on sugarcane

The classification results, collected as outlined above in Section 6.3.1, are set out in Table 6.1, and show poor accuracy ($< 66\%$) for all feature extraction and classification techniques used for discrimination of sugarcane and guinea grass. The most effective result was the LBP feature extraction technique with the SVM classifier at 65% and the worst results were from a GLCM feature extraction technique and MLP classifier at 49%. These results were not unexpected as the target plants were both grass-like. The results demonstrated that traditional textural feature extraction and classification techniques were not effective for discriminating between guinea grass and sugarcane.

Table 6.1: Classification results of existing techniques for guinea grass.

Feature extrac- tion technique	SVM accuracy	MLP accuracy	k-NN accuracy	Naive Bayes accuracy	Training set size*	Test set size*	Best window size
LBP	65.4%	50%	60.8%	57.8%	310	332	17
GLCM	57.7%	49.5%	52.6%	51.5%	198	194	27
GLRLM	62.8%	54.1%	63.3%	51.4%	218	218	27

* size is given as the number of frames

6.3.3 Results of existing classification techniques on pyrethrum

Evaluation for the identification of pyrethrum from all other plants was initiated by using the same textural feature extraction and classification techniques used for the sugarcane initial evaluation and outlined above in Section 6.3.2. The results for the LBP, GLCM and GLRLM are shown in Table 6.2. The highest positive identification rate for LBP, GLCM and GLRLM was 63% and the lowest positive identification rate was 47%. The LBP, GLCM and GLRLM results were not satisfactory and new techniques need to be developed to allow better classification.

Table 6.2: Classification results of existing techniques for pyrethrum.

Feature extraction technique	SVM accuracy	MLP accuracy	k-NN accuracy	Naive Bayes accuracy	Training set size*	Test set size*
LBP	55.3%	50%	52.6%	52.6%	202	202
GLCM	55.3%	57.9%	47.1%	50%	202	202
GLRLM	57.9%	50%	63.2%	57.9%	202	202

* size is given in number of frames

6.4 Evaluation methodology for real-world spot spray performance

6.4.1 Evaluation method

As reviewed in Chapter 2, the standard criterion used in the literature for the effectiveness of an algorithm in plant identification is the correct identification rate. However, for a real-world weed spot spray algorithm, additional criteria are required to provide an evaluation that can address the algorithm's economical

viability as well as correct identification rate. Therefore, in this research, the evaluation is assessed using three accuracy rates:

- correct target identification rate (hit rate);

- incorrect target identification rate (miss rate); and

- false positive rate (false triggers).

Hit rate is the percentage of the targeted weed sprayed with herbicide; the miss rate is the percentage of targeted weed not sprayed with herbicide; and the false trigger rate is the percentage of crop that has been misclassified as weed and sprayed with herbicide.

6.4.1.1   Acceptable accuracy rates

Ideal rates for weed identification are 100% hit rate, 0% miss rate and 0% false trigger rate. The spot spray algorithms can be adjusted (tuned) to achieve a result as close as possible to the ideal rates. In practice, the adjustment is targeted for a particular spray application (e.g. increasing hit rate at the expense of more false triggers) and will vary depending on field conditions, weed infestation, crop density and tolerance for crop losses. These variations in hit, miss and false trigger rates make it difficult to identify specific targets for the rates. Therefore, test results are compared to the ideal values of hit, miss and false trigger rates to determine the most effective algorithm and growth stage of operation.

6.5   Custom classification technique for guinea grass in sugarcane

6.5.1   Object Tracking Classification (OTC) technique for sugarcane

Tracking component position from frame to frame was identified as a potential method to enhance classification accuracy and a technique labelled Object Tracking Classification (OTC) was developed. From visual observation, it was determined that the tracking of the retained, segmented components of the image in sequential video frames was different for the shorter, denser clumped grasses (guinea grass, couch grass and nut grass) than for sugarcane.

The sugarcane components were sorted into the deleted image during segmentation by the DCSA. However, intermittently, sugarcane leaf components were retained, due to known limitations in the DCSA (Section 1.9.1) at multiple leaf points of intersection (height and colour). The shorter, denser clumps of guinea grass tended to have components in each sequential video frame that could be tracked from the top of the ROI to the bottom of the ROI over consecutive video frames as the camera moved over the plant. The OTC required that:

- a component must appear in a consecutive number of frames, greater than CF_{thresh}. Therefore the parameter CFthresh acts as a filter size, filtering out the intermittent sugarcane leaf components in the retained components image.;

- the component appears in the same columns (X axis) of the image between frames; and

- the component appears in increasing row numbers (Y axis) of the image between frames.

6.5.2   Objectives of OTC field trials

The results for the hit, miss and false trigger rates are given below in Section 6.5.4. The literature review identified that occlusion was a primary cause for error in targeted weed identification and the development of the algorithms in this research has been aimed at this problem. Therefore, occlusion has been evaluated in the results in Section 6.5.5.

6.5.3   Guinea grass size definition and OTC experimental setup

The results provided for guinea grass identification in sugarcane uses the sizes of weeds (small, medium and large) as previously defined in Section 5.6.3. For convenience the size definitions are repeated here, and in Figure 6.2:

- Small - A guinea grass plant fitting in a circle of between 0.1 to 0.2 m diameter and occupying at least 30% of the circle, by pixel area.

- Medium- A guinea grass plant fitting in a circle of between 0.2 to 0.4 m (inclusive) diameter and occupying at least 40% of the circle, by pixel area.

- Large - A guinea grass plant fitting in a circle of between 0.4 to 0.8 m diameter and occupying at least 40% of the circle, by pixel area.

Figure 6.2: Image of medium weed inside a 0.3 m circle approximately 40% filled (as Figure 5.14).

A subset of the data collected from the different growth stages was used for evaluation of the developed algorithm. The subsets of data were from four replicate runs collected on 10/10/12, field 28-B (short sugarcane); 10/10/2012, field 13-A (medium sugarcane) and 6/11/12, field 2-A (high sugarcane). This data was typical of the crop and weeds encountered on 'Fairymead' farm. The size definition of the sugarcane is; short 0.1 m - 0.49 m; medium 0.5 m - 1 m; and high 1.01 m - 1.3 m, as outlined in Section 5.3.3.

During development, it was observed that the false positive rate was affected by the threshold CF_{thresh} set for the number of successive frames the clump was identified in when tracked. The results in Tables **??** and 6.3 were obtained with a threshold CF_{thresh} setting of four determined by trial and error as optimal. If CF_{thresh} was set lower (e.g. two or three), the false trigger rate increased. If CF_{thresh} was set higher (e.g. five or six), the false positive rate decreased, but so did the correctly identified rate as outlined above (Section 6.4.1).

6.5.4   OTC field trials – results and discussion

Table 6.3 shows the hit miss and false trigger rates as well as the standard deviation of the data in the four runs for each sugarcane height.   The accuracy results for the different sugarcane growth sizes in Table 6.3 show that over the eight runs for medium and high sugarcane, the average hit rate was 90% and 83% respectively leaving an average miss rate of 10% and 17% respectively.   A standard deviation of $\pm$ 4% and $\pm$ 5% respectively for both hit and miss rate highlight the consistency of the algorithm performance.

Table 6.3: Hit, miss and false trigger rate results of object tracking classification in sugarcane.

Sugarcane size	Individual run hit rate	Individual run miss rate	Individual run false trigger rate	Total hit rate	Total miss rate	Total false trigger rate
High	89%	11%	4%	83%	17%	5%
	85%	15%	4%			
	81%	19%	5%	$\pm$5%	$\pm$5%	$\pm$1%
	76%	24%	6%			
Medium	95%	5%	1%	90%	10%	3%
	84%	16%	1%			
	91%	9%	5%	$\pm$4%	$\pm$4%	$\pm$2%
	89%	11%	4%			
Short	100%	0%	45%	85%	15%	32%
	55%	45%	16%			
	100%	0%	36%	$\pm$19%	$\pm$19%	$\pm$11%
	86%	14%	33%			

False triggers waste herbicide in spot spraying weeds in a fallow situation. A high false trigger rate is a larger problem when spot spraying weeds in a growing crop, as the false triggers kill the crop. Table 6.3 shows that the medium growth stage had a false trigger rate of 3% and high growth stage had a false trigger rate of 5% and a standard deviation in the datasets analysis of $\pm$ 2% and $\pm$ 1% respectively, highlighting consistency in the algorithm performance.

Short sugarcane had a hit rate of 85% and a miss rate of 15% which is between the results for medium and high sugarcane, but the standard deviation for the hit and miss rates in the datasets was 19% which displays inconsistent performance. The false trigger rate for short sugarcane was 32% with a standard deviation of ± 11% which indicates less consistent operation of the algorithm at this growth stage. A 32% false trigger rate in the small sugarcane would kill one third of the crop.

A reason for the difference in the accuracy rates between the short sugarcane and the medium and high sugarcane is that the discrimination capability included in the DCSA did not function as effectively on the small sugarcane. The DCSA performance was lower on the short sugarcane as there was no difference between the sugarcane and guinea grass in height, leaf size and shape. When sugarcane is in the medium to high growth stages the leaves are long and thin (e.g 0.8 m long and 0.06 m wide) and are deleted in the segmentation but when the growth stage is short, the leaves are short and thick (e.g. 0.1 m long and 0.04 m wide) and retained.

6.5.5   OTC field trials with respect to occlusion

6.5.5.1   Test setup and real-time frame rate with respect to occlusion

Results for OTC with respect to occlusion were obtained from 856 frames used in the segmentation results Section 5.6.3 recorded on the 10/10/2012, 13-A (Table 3.1), 'Fairymead' farm, Bundaberg; and a further 1140 consecutive video frames from the same recording, totalling 1,996 frames over a physical distance of 140 m. The video sequence of 1,996 contained 47 guinea grass weeds and approximately 40% of the frames contain nut grass ranging from sparse (<10 % ground covered) to complete coverage. This subset of data was chosen as it was within the optimum growth parameters for spot spraying (Section 3.4.2.2).

The video frame recording rate was 12 fps and the camera traveled 0.07 m (groundspeed) between consecutive video frames. A distance of 0.07 m between frames is a groundspeed of 8 km/h (the maximum commercial groundspeed) when run at 30 fps. A frame rate of 30 fps is achievable with the novel synchronised pipeline processing technique (Section 7.5) developed as part of this research. Table 6.4 presents the results for guinea grass identification, which used the same settings as used in segmentation for sugarcane (Section 5.6.2.1), and the parameter CF_{thresh} set at four frames, as determined above in Section 6.5.4.

6.5.5.2   OTC with respect to occlusion – results and discussion

Table 6.4 shows there was a total of 47 guinea grass plants in the data with 41 being correctly identified as guinea grass. There were six small guinea grass plants missed and by visual inspection of the images, the components segmented by the DCSA were smaller than the small definition above in Section 6.5.5.1. Of the six missed small guinea grass plants, four were occluded but were potentially larger than the small guinea grass size when the area under the occlusion is taken into account. Therefore, although the DCSA can segment the plants from each other, occlusion can still cause errors by reducing the apparent size of the plant (Figure 5.13). From the data provided in Table 6.4 the following was determined:

- A total (occluded and non-occluded) guinea grass 'hit' rate of 87%.

- A total (occluded and non-occluded) 'miss' rate of guinea grass of 13%.

- A 'miss' rate of guinea grass due to occlusion of 8%.

- A 'hit' rate of occluded guinea grass plants of 86%.

Table 6.4: Results for guinea grass identification with respect to occlusion.

Plant size	Correctly identified occluded guinea grass	Correctly identified non-occluded guinea grass	Number of guinea grass missed	Total
Small	12	9	6	27
Medium	11	5	0	16
Large	3	1	0	4
Total	26	15	6	47

Table 6.5 shows there was a total of 24 non-guinea grass weeds identified as guinea grass and that two of these were broadleaf. By inspecting the images, the nut grass and other grass that triggered a hit was in dense (>70% ground covered) low patches that had the same depth and colour criteria as the small and medium guinea grass. The known DCSA limitation of component splitting (Section 5.4.2.2) created components that, intermittently, are in contiguous frames and would be tracked by the classification algorithm, causing false triggers. The intermittent false triggers on other weeds is not a significant issue as the destruction of all weeds is advantageous to the sugarcane crop.

Table 6.5: False trigger results in guinea grass identification from sugarcane.

Nut grass	Broadleaf	Other grass	Total other weeds	Sugarcane
18	2	4	24	11

There were also 11 instances of sugarcane plants being incorrectly identified as guinea grass. A false trigger rate on sugarcane of 3.9% was determined by dividing the incorrect identification of sugarcane (11) by the total numer of sugarcane plants. The total numer of sugarcane plants was found by multiplying the plant linear density (two per metre, from Section 3.4.2.1) by the length of run (140 m). As mentioned in Section 6.5.5.1, the false trigger rate can be adjusted by changing

the threshold CF_{thresh} and this will also effect the guinea grass hit rate.

6.5.6 Identification of guinea grass in sugarcane – summary and conclusions

6.5.6.1 Existing texture feature extraction techniques

From the results in Table 6.1, the accuracy of the existing techniques in discriminating guinea grass from sugarcane in the image was unsatisfactory. When the feature extraction techniques shown in Table 6.1 were evaluated by the research cited in the literature review, they were implemented on images where the total image was grass or broadleaf but not both. In application of image analysis in multi-plant images (i.e. sugarcane and weed) the Depth Colour Segmentation Algorithm (DCSA) separated the image into smaller components which were made up of only one plant. However, smaller segmented components may not provide enough repetition in texture for texture feature extraction techniques to provide consistent results. The repetition in texture is provided when carried out on large components and/or whole images as seen in the literature review.

The accuracy results show that discrimination of plants of similar species (i.e. grass from grass type and broad leaf from broadleaf type) with texture operators is difficult. This is because the texture of these plants varies greatly within their species, unlike the Brodaz test sets (Brodatz 1966), which are commonly used to test texture extraction techniques.

6.5.6.2 OTC technique

The OTC technique results show that the overall combination of the DCSA and sequential tracking of the components through consecutive video frames is an effective means of identifying guinea grass. This technique has an effective guinea

grass identification rate of 87%, being able to identify 86% of occluded guinea grass.

The OTC technique was able to identify all guinea grass that was large or medium, and the small sized guinea grass that was above 40% of the area filled (Figure 6.2). Therefore, the effective minimum size of guinea grass targeted by the technique, given for the calibration settings used is a guinea grass plant fitting in a circle of between 0.1 m to 0.2 m diameter and filling the circle by 30% of its area or approximately 0.05 m^2.

The false trigger rate was 3.9%, with the threshold $CF_{\text{thresh}} = 4$. However this should be looked at with respect to the classification settings in the DCSA as well as CF_{thresh}. By changing the classification settings in the DCSA, and the value of CF_{thresh}, smaller or larger guinea grass plants can be targeted as discussed in Section 6.5.5.1. However, changing the settings will also affect the false trigger rate on sugarcane. The false trigger rate on non-sugarcane plants is not considered a problem, as spraying and killing non-sugarcane plants will enhance the sugarcane crop. Ultimately, settings for the DCSA and CF_{thresh} are a commercial decision to be made by the farmer when setting up the spot spray system in the field (Section 6.4).

The processing time required for OTC of guinea grass is addressed by the novel synchronised pipeline processing technique developed in this thesis for real-time spot spraying (Chapter 7).

6.6 Custom classification algorithms for weeds in pyrethrum

6.6.1 Development of pyrethrum classification algorithms

Visual inspection of the pyrethrum data identified that height and spatial position may improve the outcome of the classification result. Height was a reliable

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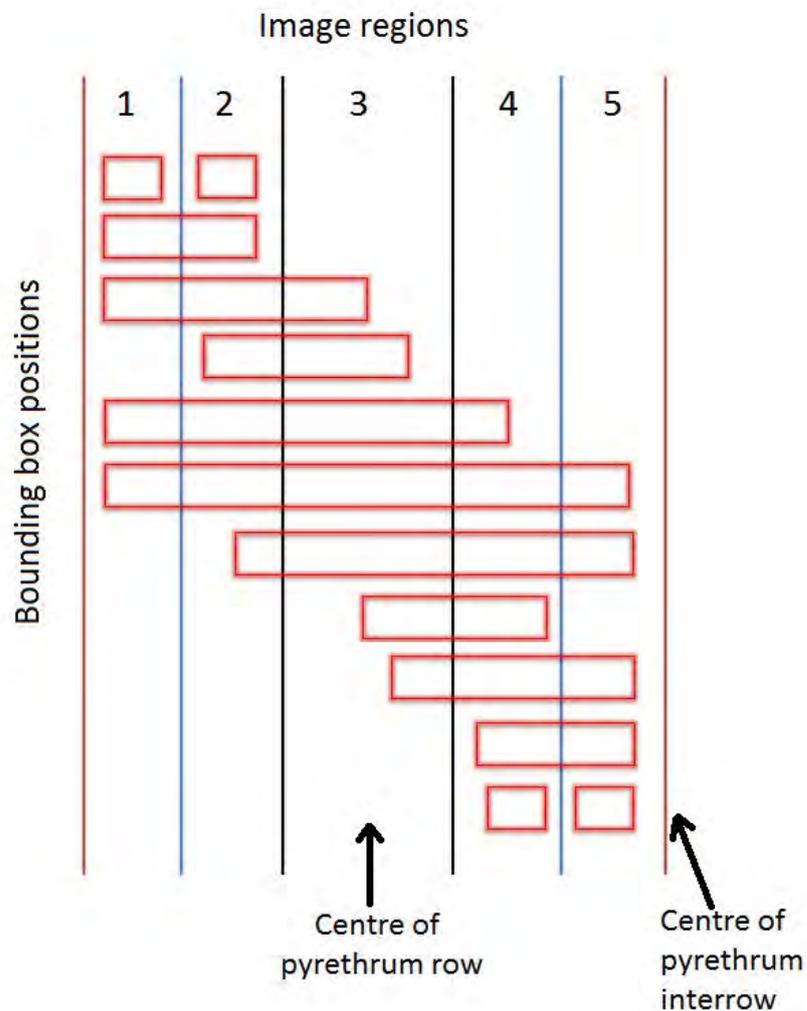
Figure 6.3: Bounding box positions used in spatial analysis that determine if a component is weed.

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.
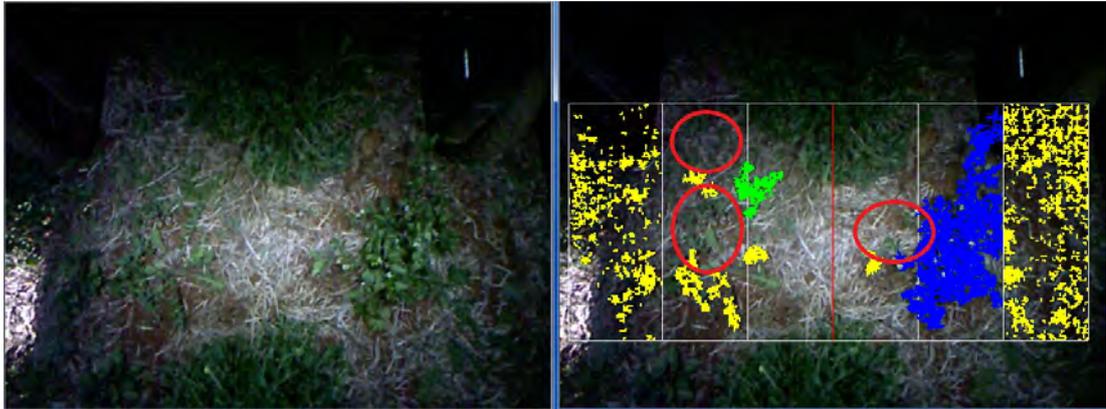
Figure 6.4: Results of the spatial segmentation method on small weeds highlighting misses from poor binarisation in the red ellipses. RGB image of pyrethrum and weeds on the left hand side. The images on the right hand side shows the weeds identified by shading the pixels were yellow, green, red, and blue depending on component size and position.
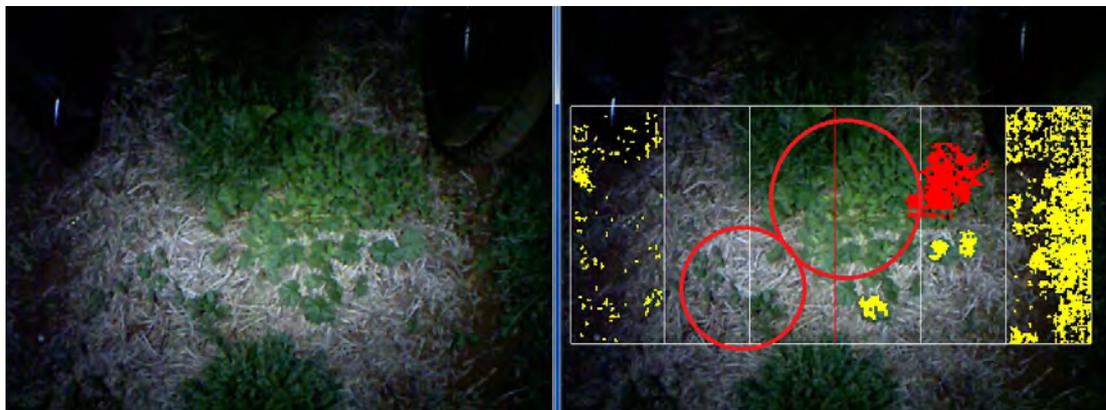
Figure 6.5: Results of the spatial segmentation method on large weeds highlighting misses from occlusion in the red ellipses. RGB image of pyrethrum and weeds on the left hand side. The images on the right hand side shows the weeds identified by shading the pixels yellow, red, blue and green depending 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 $minWeed_{\mathrm{thresh}}$, 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.
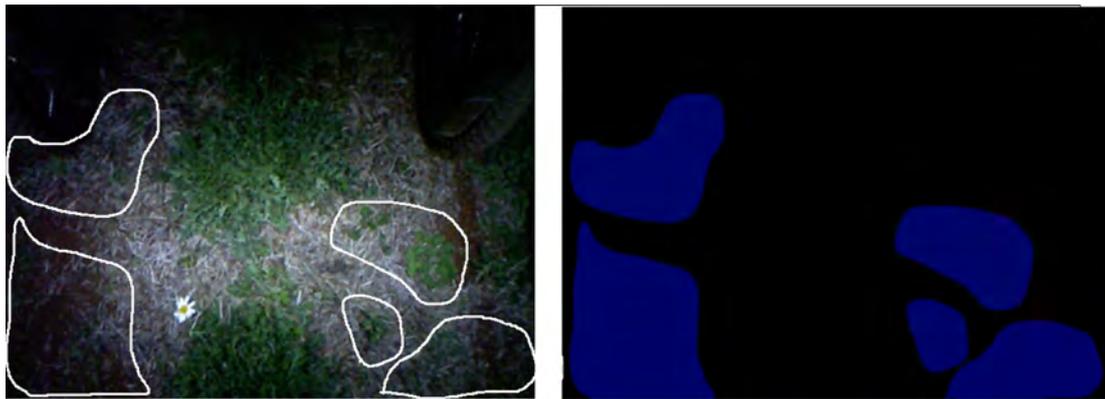
Figure 6.6: Images showing the ground truthing mask for the evaluation software. The image on the left shows the weeds circled by a mouse in white. The image on the right is the mask with the circled areas of the first image highlighted in blue.

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.

Chapter 7

Real-time processing

Automated weed spot spraying requires the spot application of chemical to a weed whilst traveling at speeds up to 10 km/h in the row crop, horticulture, and sugarcane industries in Australia, and up to 20 km/h in the Australian broadacre industry. To achieve 'on-the-go' capability an automated weed spot sprayer needs to operate in real-time.

7.1   Chapter outline

This chapter contains six sections which provide information on real-time computing and its implementation into an 'on-the-go' automated spot sprayer.

- Section 7.2 provides a general overview of what a real-time system is in relation to computing systems.

- Time constraints/limitations are a significant factor in a real-time system, and Section 7.3 details these limitations and the interaction between frame rate, groundspeed, and analysis time. This section addresses the requirements in a real-time automated spot spraying system, and the consequences

if real-time functionality is not adequate, such that weeds will be missed by
the sprayer.

- Real-time systems can be implemented in CPU-based computers and hard-
  ware systems such as Programmable Logic Devices (PLD). Section 7.3.3
  discusses the implementation onto a PLD.

- CPU based implementations are set out in Section 7.4 along with their pro-
  gramming techniques. This section discusses implementations onto CPUs
  operating with traditional single-core techniques and also parallel process-
  ing on multicore CPUs. It outlines the history of both sequential and asyn-
  chronous operating techniques and shows that the asynchronous parallel
  programming techniques cannot always guarantee real-time deadlines, also
  why 'just adding more cores' does not provide a linear speed increase.

- The aim of the research undertaken was to develop a real-time automated
  spot spray system and Section 7.5 provides a detailed description of a pro-
  cessing technique developed in this research that can achieve this aim.

- Section 7.6 compares the timing results of the technique developed in this
  research to traditional single-core, and parallel processing techniques, show-
  ing that the new technique can undertake significantly more computation.

- Finally a discussion of the significance of the new technique to the auto-
  mated spot spray industry is outlined in Section 7.7.

7.2   Real-time systems overview

7.2.1   Real-time systems definition

The concept of real-time computer systems is not new. Lin & Burke (1992) stated
that a "real-time computer must produce a correct result within a specified time".

Stankovic (1992) reported that "Real-time systems are defined as those systems in which the correctness of the system depends not only on the logical result of computation but also on the time at which the results are produced". West (2001) defined a real-time system as a system where the "the results will be provided when they are needed". In this thesis, real-time is 33 ms per frame for 15 km/h groundspeeds (Section 7.3.2).

7.2.2 Real-time computation requirements

Halbwachs (1993) states the principal attributes of a real-time system, in terms of process implementation, must have the following features:

- **Concurrency**. There must concurrency between the system and the operational environment. The image analysis within an automated spot spray system for weeds must maintain a computational output that is synchronised with the groundspeed of the system.

- **Strict time requirements**. There must be strict time requirements for the system output response relative to the system input. The image analysis must be performed, a result determined, and the spray nozzle activated to apply the herbicide to the weed, while the system is positioned over the weed. If the analysis overruns the strict time requirement the spot spray system will have traveled over the weed before a result is determined, and the weed will be missed.

- **Deterministic**. The output is determined by input data and input time/s. The maximum analysis run time must be known and operate within the input frame rate of the system.

- **Reliability**. The software system implementation must be robust. Software crashes are not acceptable.

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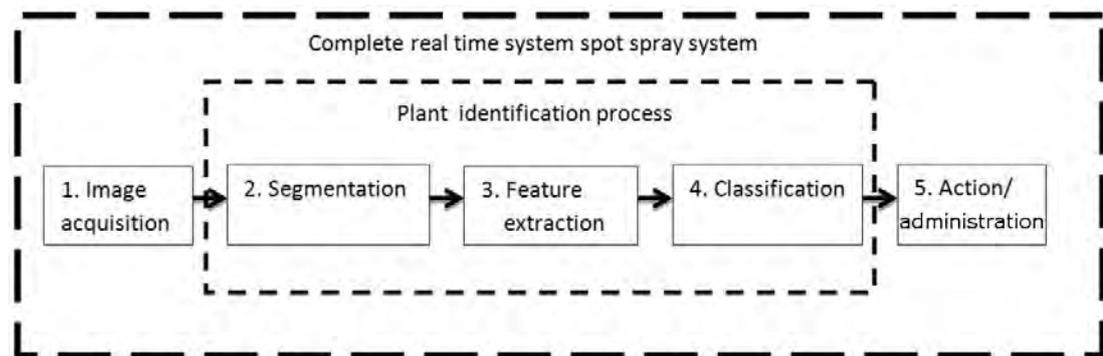
Figure 7.1: Reprint of block diagram of the spot spraying system based on real-time machine vision.

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.
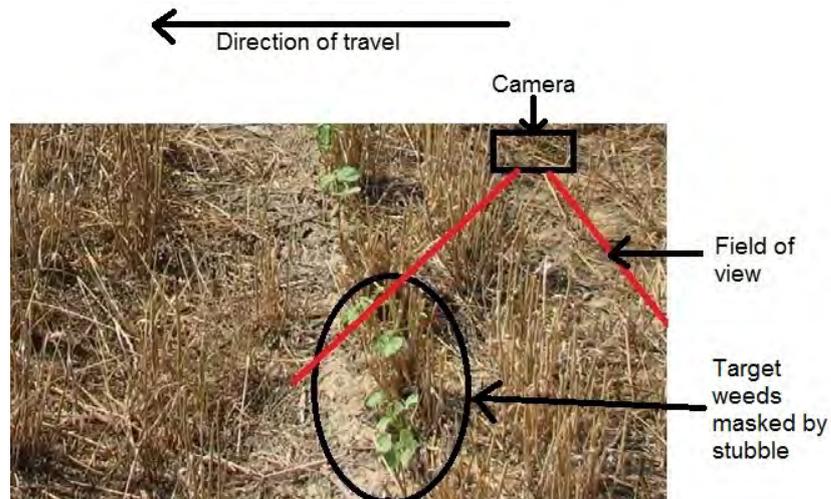
(a) Camera positioned right of weeds, occluded by wheat stubble.
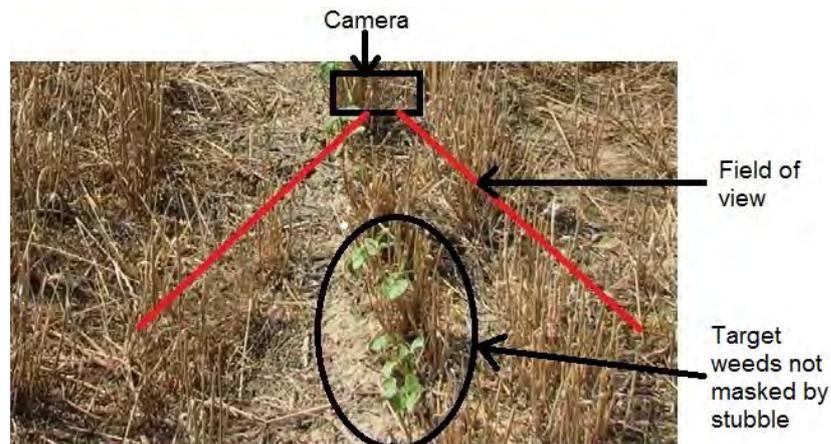
(b) Camera positioned above weeds and not occluded by wheat stubble.
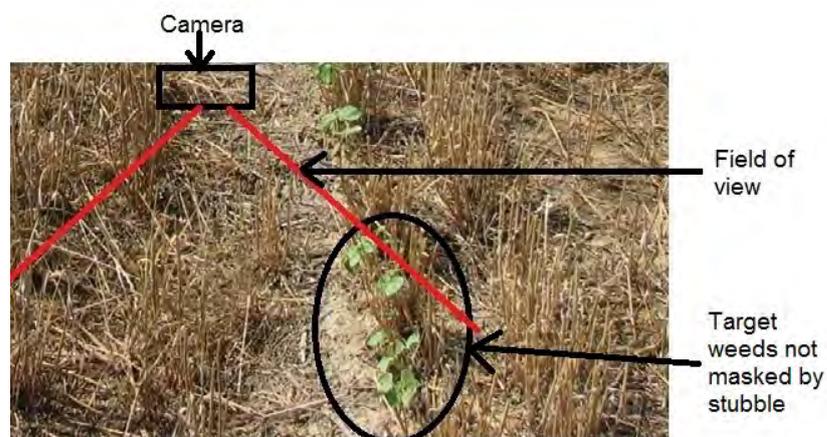
(c) Camera positioned left of weeds, not occluded by wheat stubble.

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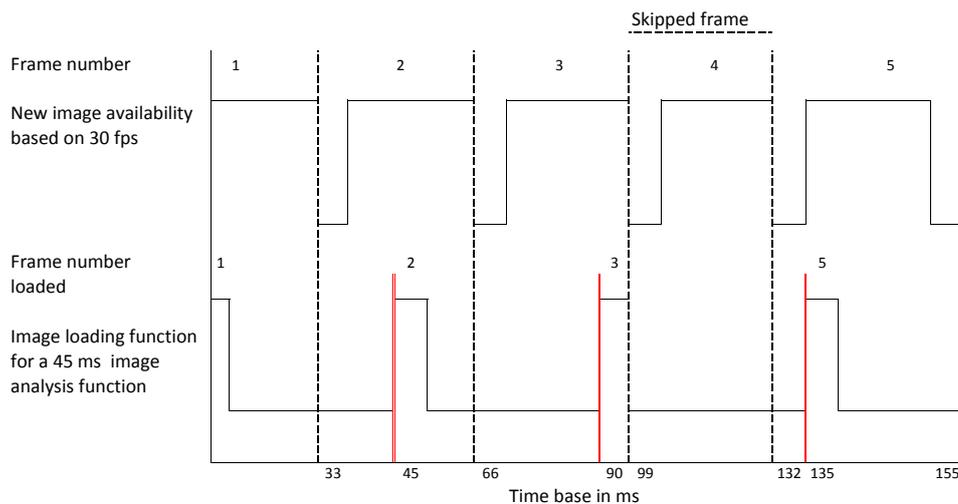
Figure 7.4: Timing digram showing an image analysis system missing an input frame of data due to the image analysis overrunning the image capture interval.

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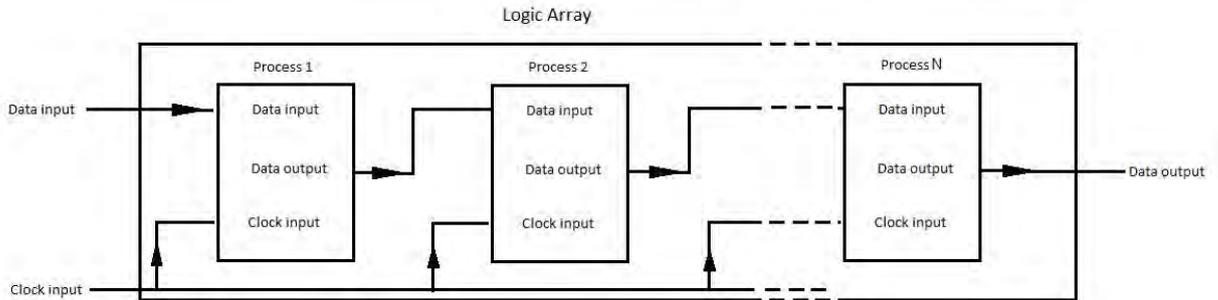
Figure 7.5: Block logic array illustrating data flow through an N-process hardware logic pipeline controlled by synchronous clocking.

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:

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

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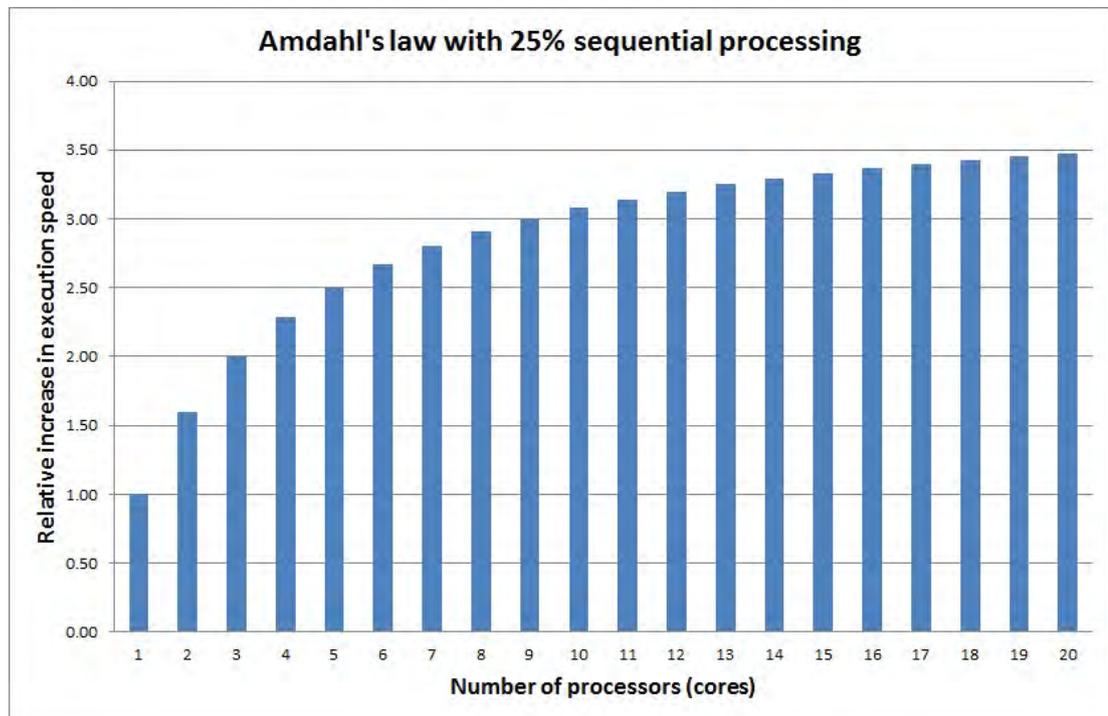
Figure 7.6: Graphical representation of the relative execution speed improvement determined by Amdahl's law with 25% sequential processing over 20 cores.

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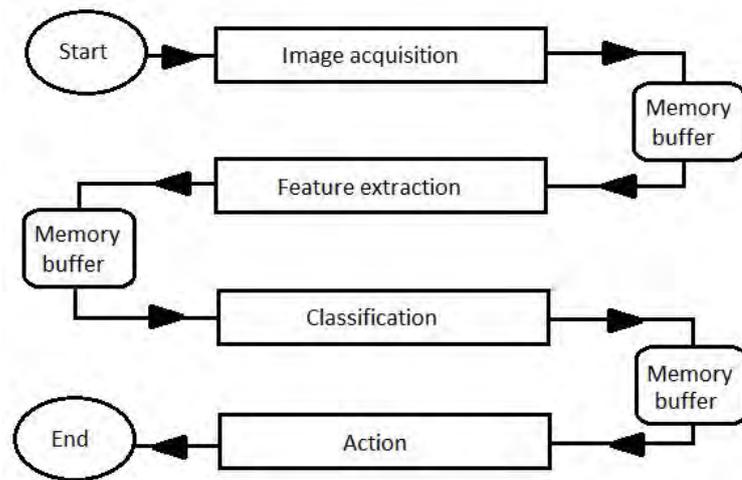
Figure 7.7: Asynchronous parallel pipeline process applied to image analysis
(adapted from (Campbell & Miller 2010)).

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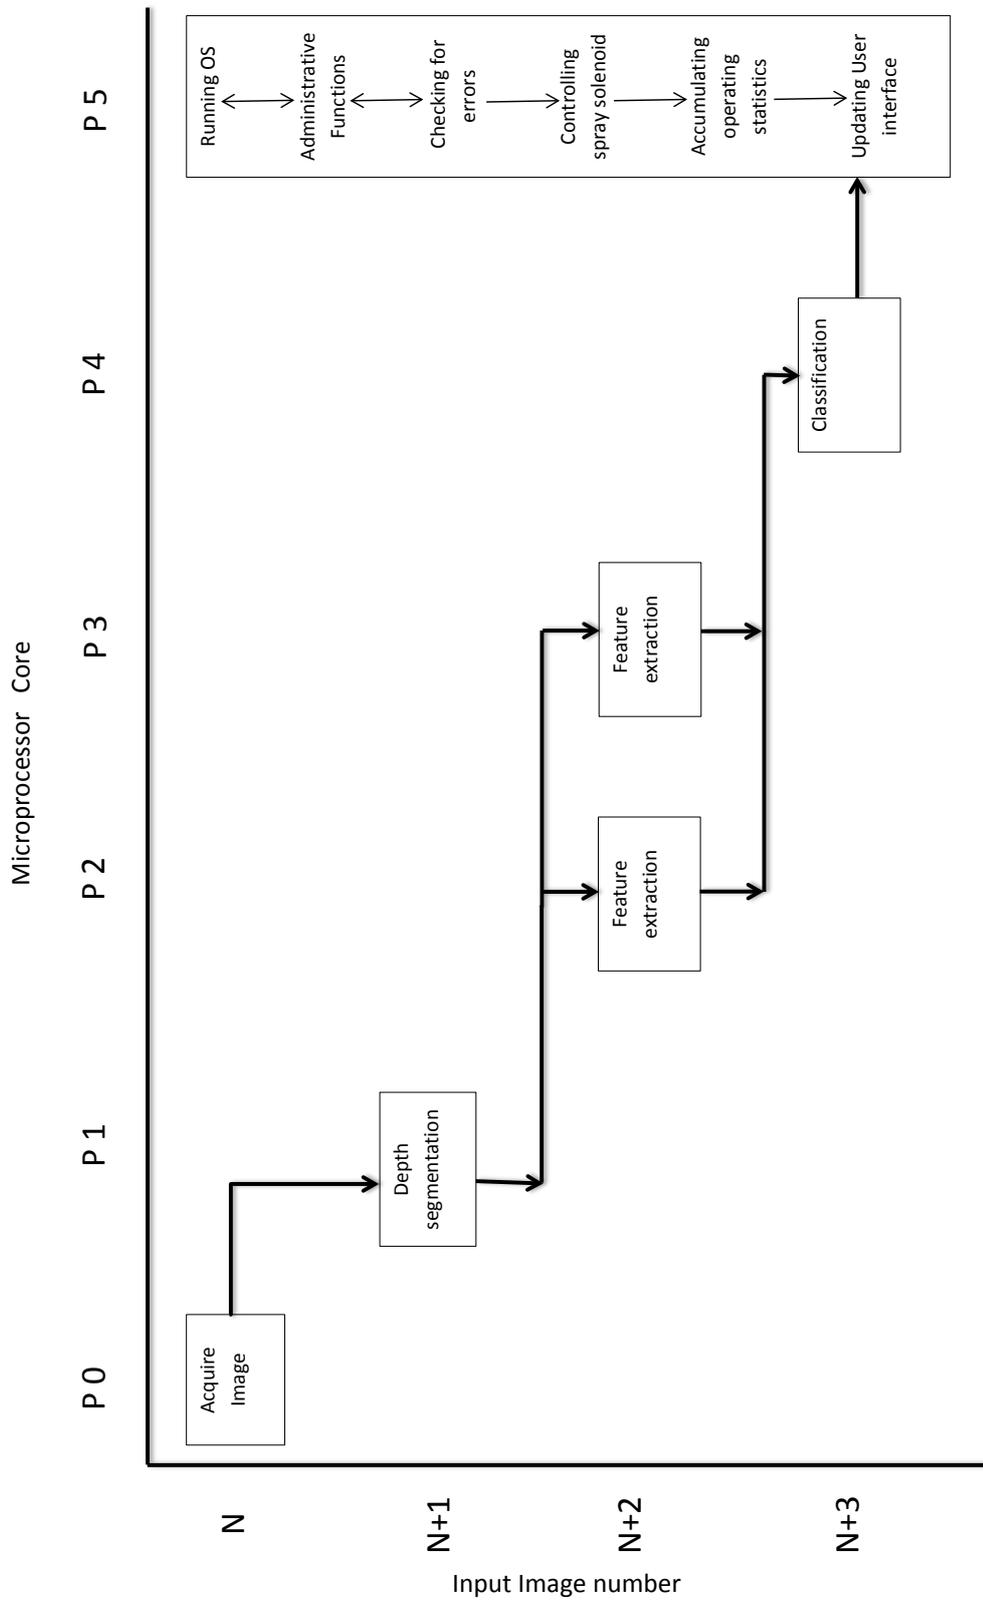
Figure 7.8: Flow diagram of the Synchronised Pipeline Processing (SPP) technique, indicating the principal processing tasks in its application to spot spraying.

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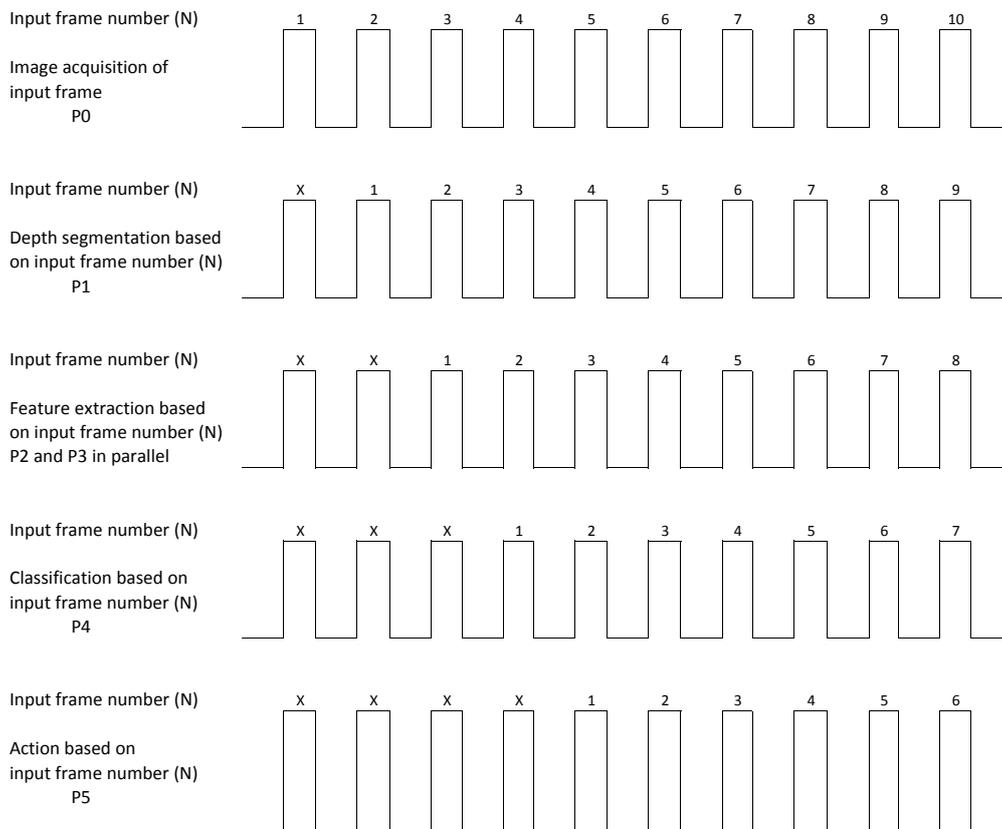
Figure 7.9: Timing diagram of the novel SPP technique corresponding to
Figure 7.8. X denotes an unknown state of the system.

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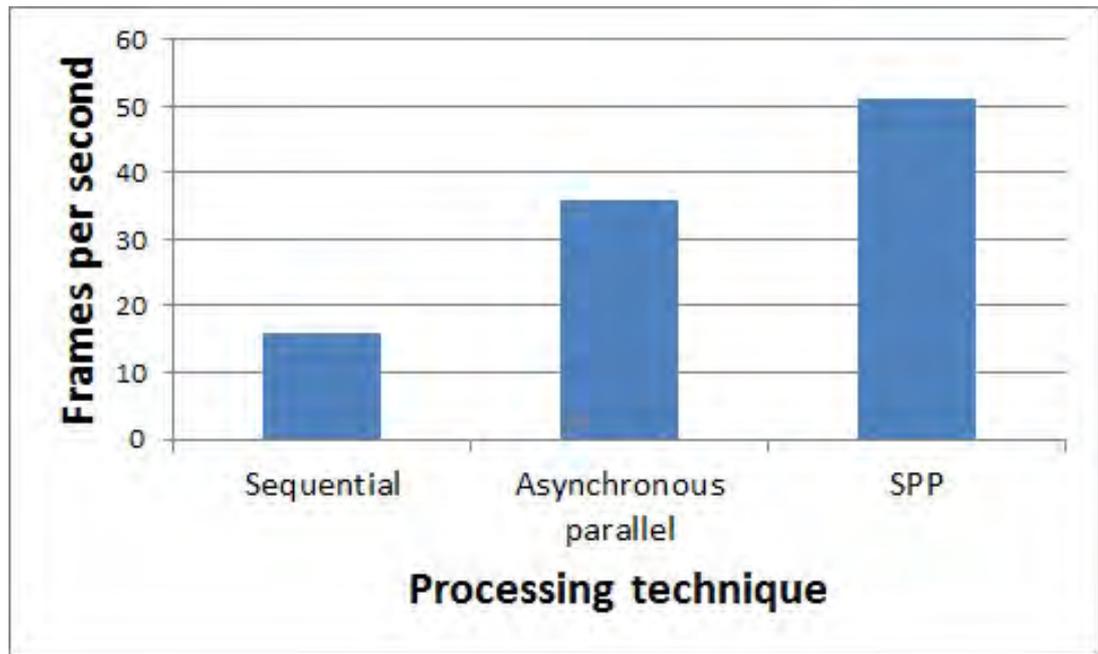
Figure 7.10: Overall improvement in frame rate of processing techniques for
sugarcane.

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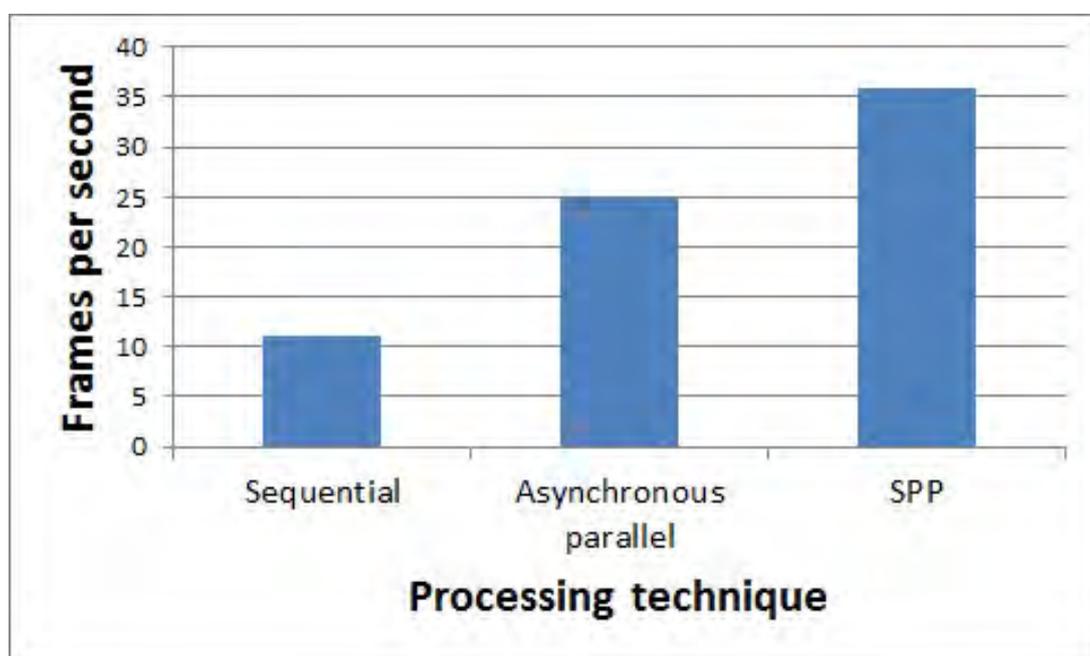
Figure 7.11: Overall improvement in fps of compared techniques for pyrethrum.

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 soft-

ware 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

occlusion tolerance and data reduction that was used in Chapter 5 for segmentation, is used in this chapter to compare the DCSA's operation in sorghum and mung beans with sugarcane and pyrethrum respectively. The evaluation method is detailed below in Sections 8.3.1 and 8.4.1.

8.2 Data collection

Data was collected on the Darling Downs in Queensland using the pyrethrum single crop-row data gathering device outlined in Section 3.5. The single crop-row data gathering device had a groundspeed of 3 km/h but fluctuated ±1 km/h depending upon the field conditions, i.e the rougher the ground, the slower the groundspeed.

The mung bean data was collected from 'Wolonga' (latitude -27.594775 , longitude 151.296262) on the 31st of January 2014 at 1:30 pm and consisted of 150 frames of colour and depth video. The sorghum data was collected from 'Kurralinden' (latitude -27.594318, longitude 151.291176) on the 13th of November 2013 at 4 pm, and consisted of 150 frames of colour and depth video.

8.3 Application of the DCSA to sorghum

8.3.1 Weeds from sorghum

The crop of sorghum was required to be discriminated from weeds. Sorghum is a grain crop which grows during the summer months in north eastern Australia. Sorghum is commonly planted in 1 m rows by a precision planter, but is sometimes grown on narrower rows planted by broadacre air seeders set up for planting wheat (which is a common crop in broadacre farming in Australia). Sorghum can grow

up to 1.5 m high depending upon the variety and growing conditions. From visual inspection of colour and depth data from the data collection device, the weeds were observed not to be grass-like and the majority of the weeds were lower in height than sorghum. Sorghum grows with grass-like leaves from a main stem (Figure 8.1.

Figure 8.1: Sorghum planted on 1 m row spacing at the 50 cm growth stage. Taken at 'Kurralinden' on the 13th of November 2013.

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). Therefore, modifications were made to the DCSA analysis software described in Section 5.5.2, namely:

- binarised segmentation technique (BST) was applied to the colour image;

- recording both retained and deleted components; and

- manual assessment of occlusion in the image before and after segmentation.

The recorded information allowed the DCSA to be compared to the BST image. A depth-only segmentation technique was not evaluated for the same reason as set out in Section 5.2.4, namely the problems associated with the depth-only connected component technique incorporating the ground in the components.

For the evaluation, the components in the 'retained image' after application of the DCSA are those components:

- whose average height was above height1 (*H1*=45);

- are green; and

- are above a threshold size of 100 pixels.

The analysis was undertaken by applying the DCSA and binarised segmentation technique ($G > R$ and $G > B$) to a video stream of sorghum and weeds with the sorghum at a growth stage of 50 cm. In the image series Figures 8.2 to 8.5, the following image analysis steps are presented:

- A colour image of sorghum (Figure 8.2).

- A BST image of Figure 8.2 (Figure 8.3).

- Components determined by DCSA applied to Figure 8.2 (Figure 8.4).

- The original image (Figure 8.2) with components retained by the DCSA highlighted and overlaid in yellow (Figure 8.5).

This sequence of images pictorially shows the DCSA process for sorghum and its sorting capability.

Figure 8.2: Colour image of sorghum plant with weeds.
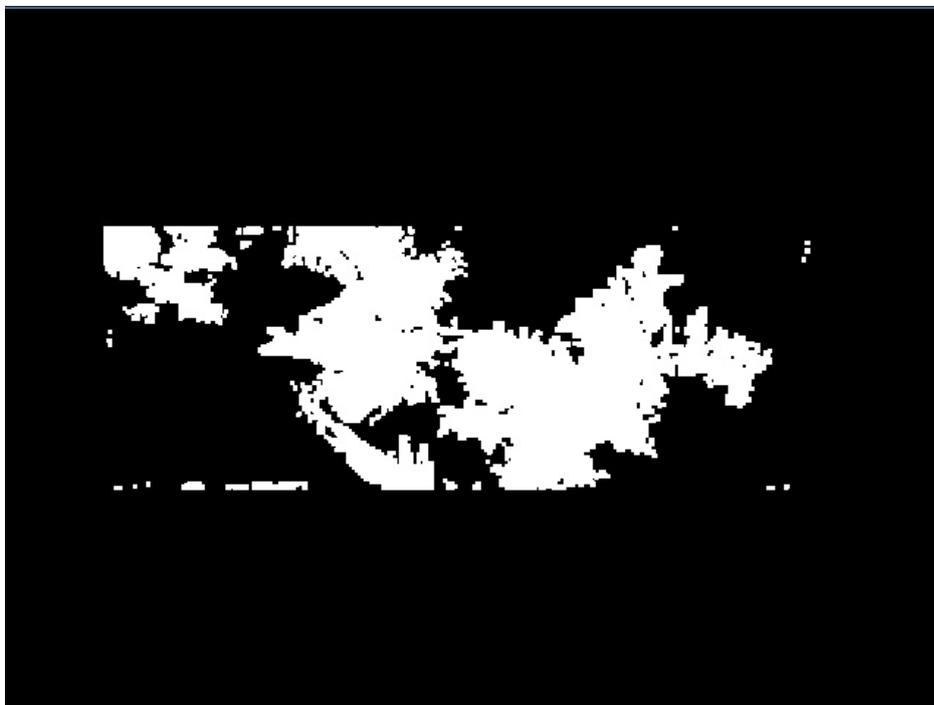
Figure 8.3: BST segmented image of Figure 8.2.

Figure 8.4: All components identified in the DCSA (in false colour) of the colour image of sorghum, Figure 8.2.

Figure 8.5: Colour image of sorghum plant and weeds overlaid with the retained components highlighted yellow.

8.3.2   Results for the DCSA used in sorghum

Table 8.1 displays the results for occlusion tolerance for the DCSA and the binarised segmentation technique application to sorghum and weeds. Table 8.1 shows there were 37 weeds in the 150 consecutive frames of video of which 23 weeds were occluded by sorghum. The binarised segmentation technique was not able to isolate any of the occluded weeds, with 100% of occluded weeds being grouped with sorghum components. However, the DCSA grouped 0% of occluded weeds with sorghum.

Table 8.1: Sorghum occlusion results from 150 frames.

Total number of weeds	37
Number of weeds occluded in video data	23
Number and percentage of weeds occluded after BST	23 (100%)
Number and percentage of weeds occluded after DCSA	0 (0%)

The sorting capability results are set out in Tables 8.2 and 8.3. From Table 8.2, there were 36 components in each frame on average and the sorting function in the DCSA removed 30 components, leaving 6 components which is a reduction in components requiring further analysis of 85%. There was a reduction in pixels requiring further analysis (Table 8.3) of 65%. The computational time of the analysis was 24 ms, which is similar to sugarcane at 22 ms. The standard deviation of the components and pixels in 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 8.2: Sorghum segmentation results for components and timing.

Statistic	Average	std deviation
Number of depth components per image	39	9
Number of depth components after segmentation	6	3
Reduction in the number of depth components	85%	-
Computation time of DCSA in ms	24	4

Table 8.3: Sorghum segmentation results for pixels.

Statistic	Average	std deviation
Depth component pixels per frame before segmentation	33584	7
Depth component pixels per frame after segmentation	11900	5278
Percentage reduction in depth component pixels	65%	-

8.4 Application of the DCSA to mung bean

8.4.1 Discrimination of weed from mung bean

The crop of mung bean was required to be discriminated from all weeds. As with sorghum, mung bean is a grain crop which grows in summer in Australia. Mung beans are grown in north eastern Australia and are commonly grown in 1 m rows planted by a precision planter however, as with sorghum it has also been grown on narrower rows planted by broadacre air seeders. Mung beans can grow up to 0.5 m high depending upon the variety and growing conditions and commonly spill out between the row (Figure 8.6). From visual inspection of the colour and depth data, it could be seen that the weeds were not grass-like and that the majority of the weeds were lower in height and vine-like. Figure 8.6 shows the crop of mung beans from which the data was collected.

Figure 8.6: Colour image of mung beans planted on 1 m rows with weeds. Taken at 'Wolonga' on the 31st of January 2014.

The evaluation software modified for sorghum evaluation was used in the mung bean data evaluation. The settings used in the DCSA for mung beans were:

- whose average height was above height1 ($H1=35$);

- are green; and

- are above a threshold size of 100 pixels.

The analysis outlined previously for sorghum was applied to mung beans at 35 cm growth stage. An example in images is as follows. The colour image is shown in Figure 8.7. The colour image (Figure 8.7) is segmented by a BST and shown in Figure 8.8. An image of the components found in the DCSA of Figure 8.7 is shown in Figure 8.9. The retained components are superimposed as yellow on the colour image Figure 8.10 to highlight the reduction in components achieved by the sorting capabilities of the DCSA.

Figure 8.7: Colour image of mung beans plant with weeds.

Figure 8.8: BST segmented image of beans plant with weeds Figure 8.7.
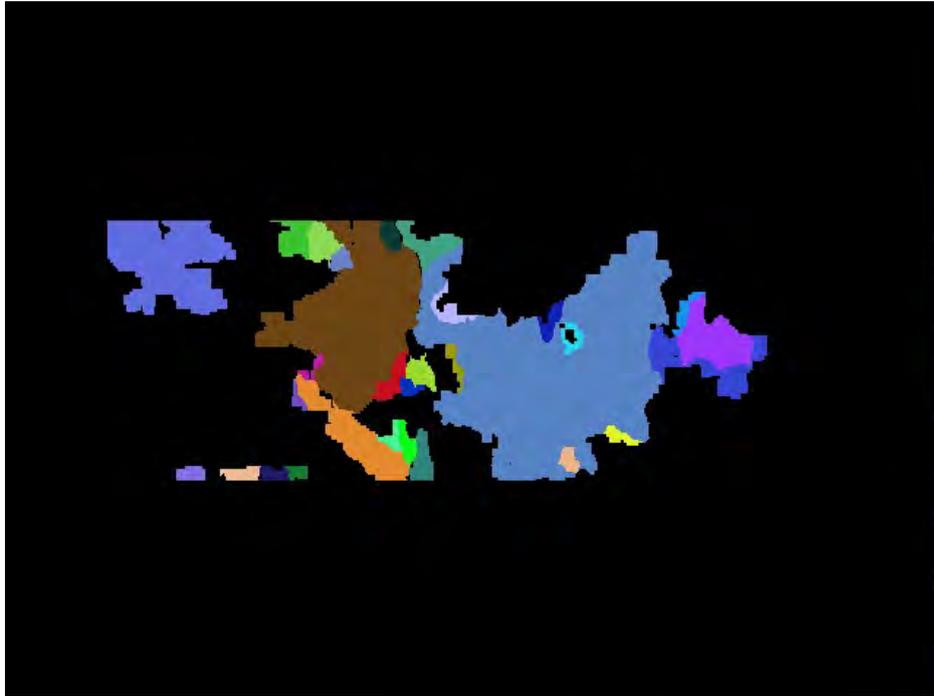
Figure 8.9: All components identified in the DCSA (in false colour) of the colour image of Figure 8.7.

Figure 8.10: Colour image of sorghum plant and weeds overlaid with the retained components highlighted yellow.

8.4.2   Results for the DCSA technique in mung beans

Table 8.4 are the results for occlusion tolerance for the DCSA and the binarised segmentation technique when applied to mung bean and weeds. Table 8.4 shows that from the 150 frames analysed there was 14 weeds in total with 6 of the weeds being occluded. The BST was not able to isolate any of the occluded weeds and the DCSA was able to segment all of the occluded weeds.

Table 8.4: Occlusion results of segmentation analysis in mung bean from 150 frames.

Total number of weeds	14
Number of weeds occluded in video data	6
Number and percentage of weeds occluded after BST	6 (100%)
Number and percentage of weeds occluded after DCSA	0 (0%)

Tables 8.5 and 8.6 display the sorting capability of the DCSA. Table 8.5 shows the DCSA removed 33 components from a total of 38 components which is an 87% reduction in the number of components requiring further analysis. Table 8.6 shows a reduction in pixels from 34884 pixels to 14329 pixels (59%) requiring further analysis. The standard deviation of the components and pixels in 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 8.5: Mung bean segmentation results regards components and timing.

Statistic	Average	std deviation
Average number of depth components per image	38	6
Number of depth components after segmentation	5	2
Percentage reduction in the number of depth components	87%	-
Computation time of DCSA in ms	25	4

Table 8.6: Mung beans segmentation results regards pixels.

Statistic	Average	std deviation
Depth component pixels per frame before segmentation	34884	5809
Depth component pixels per frame after segmentation	14329	2661
Percentage reduction in depth component pixels	59%	-

8.5 Discussion of results

The evaluation on sorghum and mung bean crops was preliminary and based on a small data set of 150 frames for each crop. However, the results from both sorghum and mung beans show a high level of occlusion tolerance for the DCSA with 100% tolerance for both data sets. The sorting capability of the DCSA was high with results for reducing pixels for further analysis of 85% for sorghum and 87% for mung beans. From visual inspection of the depth and colour data it was anticipated that the sorting results would be significant because there was a height difference between the majority of the weeds and the crop. This supports the results of Section 5.8 which shows that the sorting capability of the DCSA is higher when there is a height difference between crop and weed. The execution times would operate with the SPP at 24 ms (sorghum) and 25 ms (mung beans) which are similar to sugarcane at 22 ms. These results are judged to be promising and suggest that both the Depth Colour Segmentation Algorithm (DCSA) and the Synchronised Parallel Processing (SPP) technique are portable and may be directly applicable in other crops in the Australian broadacre and row crop industries.

Chapter 9

Conclusion and further research

9.1 Conclusion

This thesis has detailed the development and evaluation of a real-time, real-world machine vision system for automatic weed spot spraying and mapping. The developed system can discriminate crop from weed to allow the application of herbicide to the weed only.

The objectives of this thesis are addressed in Sections 9.1.1 to 9.1.4 below.

9.1.1 Objective 1: Develop algorithms incorporating 2D and 3D data

To develop an algorithm/s that incorporates 2D (colour) and depth data from video streams to achieve weed discrimination from crop in a 'real-time', 'real-world' environment at commercially realistic groundspeeds.

9.1.1.1   Develop 3D algorithms for real-world conditions

A unique segmentation algorithm called the Depth Colour Segmentation Algorithm (DCSA) was developed in Section 5.3 which incorporates 2D colour and depth data to segment occluding leaves in a real-time, real-world situation. A further feature developed into the DCSA was sorting of components to reduce the amount of data requiring feature extraction and classification. The sorting was based on colour and depth statistics of the component.

Unique feature extraction and classification techniques were also developed. An object tracking classification technique (Section 6.5.1) was developed and evaluated to identify guinea grass from sugarcane. Techniques developed for pyrethrum used depth, colour, size, spatial and textural features (Section 6.6) to identify pyrethrum from weeds

9.1.1.2   Real-time development

Real-time capability was obtained through the development of a new processing technique called Synchronised Parallel Processing (SPP) which combines the benefits of hardware processing on logic devices and consumer multi-core CPUs. The SPP technique is unique and maintains a high frame rate (which dictates the maximum groundspeed) whilst expanding the processing time by allocating the workload in a permanently allocated pipeline synchronised by the incoming video image.

9.1.1.3   Real-world data

The data in this research was collected at commercial operating groundspeeds for sugarcane and pyrethrum from data acquisition units in typical field conditions outlined in Chapter 3. The sugarcane data ranged from 0.25 m to 1.3 m in height and the pyrethrum ranged from 0.1 m to 0.5 m in width and height up to 0.35 m.

9.1.2 Objective 2: Demonstrate 3D techniques in real-time, real-world conditions

To demonstrate that the addition of depth data to a suitable image analysis technique can achieve weed discrimination from crop in a commercially acceptable operational window, i.e. at a range of crop growth stages in addition to a 'real-time', 'real-world environment.

9.1.2.1 3D technique performance

The DCSA was shown to segment plant from stubble and potentially other foreign objects even when they are a similar height to the plant material, with an accuracy when occluded of greater than 99%. Additionally the DCSA sorting produced an 84% reduction in pixels requiring feature extraction and classification in sugarcane and 55% in pyrethrum. The DCSA worked best in sugarcane at the medium (0.5 m) to high (1.3 m) growth stages and pyrethrum was best when the diameter of the plant was greater that 0.15 m.

The object tracking classification technique (Section 6.5.1) could discriminate guinea grass from sugarcane with a guinea grass identification rate of 87% (where 86% of the guinea grass was occluded) with a minimum weed size of 0.05 m^2 and a false trigger rate of 3.5%. Techniques for pyrethrum (Section 6.6) could identify pyrethrum from weeds with a pyrethrum identification rate of up to 98% and a false positive rate of 1.2% of pyrethrum plants of 0.15 m diameter and larger.

9.1.2.2 Real-time performance

The SPP was able to achieve frame rates for sugarcane and pyrethrum of 51 fps and 36 fps respectively or a top groundspeed of 18.5 km/h and 17.2 km/h respectively.

9.1.3   Objective 3: 3D versus 2D

To evaluate the performance of the developed technique under a range of real-world environment conditions; in particular with respect to 3D space versus 2D or depth on their own.

The DCSA was able to segment occluded leaves with 99% accuracy where the BST was not able to segment occlusions at all. Section 6.2 evaluated the use of typical 2D texture feature extraction techniques of Grey Level Co-occurrence Matrix (GLCM), Grey Level Run Length Matrix (GLRLM) and Local Binary Patterns (LBP) in sugarcane (accuracies between 49% and 66%) and pyrethrum (accuracies between 47% and 63%). Section 6.6 incorporated depth data with LBP data for pyrethrum identification to improve the classification taking the accuracy form 63% to 90% which is an improvement of 37%.

9.1.4   Objective 4: Portability of algorithms with respect to other crops

To demonstrate that the algorithms are adaptable to a range of crops under practical commercial conditions.

The DCSA was developed as a generic technique, was portable (Chapter 8) and was applied to four crops: sugarcane, pyrethrum, sorghum and mung beans in this thesis.

Portability of the DCSA was addressed in Chapter 8 by applying the DCSA to sorghum and mung bean crops with positive preliminary results. The DCSA was able to segment all occluding leaves and the sorting capability associated with the DCSA reduced the data requiring further analysis by 85% for sorghum and 87% for mung beans. The execution time of the DCSA was 24 ms for sorghum

and 25 ms for mung beans which is real-time capable.

9.2   Potential further research

The scope of this thesis has covered the machine vision process associated with machine vision spot spraying. The research in this thesis has identified five areas for further research which are:

1. Comprehensive evaluation of DCSA. The portability of the DCSA underwent preliminary evaluation on sorghum and mung beans as a part of this thesis. However, a full evaluation on crops grown in the Australian no-till and minimum-till cropping systems would provide increased confidence for segmentation of plants with occluding leaves in a real-world environment and encourage feature extraction and classification techniques to be widely implemented. Clearly this will involve extensive routine trials

2. A robust daylight compensation technique to allow the sensor equipment to operate on open boomsprays, i.e. not under light restricting covers. The compensation technique will need to provide shadow removal algorithms as well as white balance compensation. Daylight compensation is required for machine vision spot spray systems to become more mainstream and have greater operational diversity

3. A depth sensing device capable of direct sunlight operation at frame rates that allow groundspeeds up to 20 km/h. Typical stereo-vision cameras have difficulty in varying light conditions and require significant time (with respect to a real-time system) for correspondence matching. Structured light systems offer a cheap and effective method of depth sensing, however, they are affected by dust, which is prevalent in agricultural circumstances. A system needs to be developed that can be largely unaffected by dust,

provide accurate depth information, at a economical cost and suit a real-time situation.

4. More complex algorithms that can use the additional processing time to identify a wider weed spectrum, or improved robustness of existing algorithms. The SPP will enhance the machine vision systems to operate several feature extraction algorithms or highly complex individual algorithms. It is envisioned that additional features may enable more robust classification of weeds across a wider weed/crop spectrum.

5. Develop precise herbicide application spray technology. Spraying weeds in a crop situation requires the smallest possible herbicide footprint. Any overspray of the weed can potentially cause death to a crop plants. Spot spray delivery technology needs to be developed that provides an optimum spray pattern for standing weeds, as well as prone weeds, with little to no overspray at speeds up to 20 km/h. This project will require developing a spray manifold with numerous spray solenoids and each nozzle applying a small spray width on the ground (e.g. 5 cm, therefore requiring 20 per meter). The solenoids will need to be highly efficient electrically, with fast turn on/off times. The spray manifold will need to consider the effects of continually changing fluid pressure in the manifold due to the activation of many small solenoids.

References

Ahmad, I., Siddiqi, M. H., Fatima, I., Lee, S. & Lee, Y.-K. (2011), Weed classification based on haar wavelet transform via k-nearest neighbor (k-nn) for real-time automatic sprayer control system, *in* 'Fifth International Conference on Ubiquitous Information Management and Communication', Vol. 17, ICUIMC'11.

Ahmed, F., Bari, H., Shihavuddin, A., Al-Mamun, H. A. & Kwan, P. (2011), A study on local binary pattern for automated weed classification using template matching and support vector machine, *in* '12th IEEE International Symposium on Computational Intelligence and Informatics', IEEE Computer Society.

Akhter, S. & Roberts, J. (2006), *Multi-Core Programming – Increasing Performance through Software Multi-threading*, Intel Press, Intel Corporation, 2111 NE 25th Avenue, JF3-330, Hillsboro, OR 97124-5961.

Amdahl, G. (1967), Validity of the single processor approach to achieving large scale computing capabilities, *in* 'AFIPS spring joint computer conference Atlantic City, N.J., Apr. 1820', Vol. 30 pp. 483-485, American Federation of Information Processing Societies.

Andersen, H., Reng, L. & Kirk, K. (2005), 'Geometric plant properties by relaxed stereo vision using simulated annealing', *Computers and Electronics in Agriculture* **49**, 219–232.

Anderson, J., Hoffman, S., Shifman, J. & Williams, R. (1962), Rd825 - a multiple-computer system for command and control, *in* 'Proceedings of the December 4-6, 1962, Fall Joint Computer Conference', Vol. 2, ACM, pp. 86–96.

Astrand, B. & Baerveldt, A. (2002), 'An agricultural mobile robot with vision-based perception for mechanical weed control', *Autonomous Robots* **13**, 21–35.

Australian Weed Management (2004), 'Introductory weed management manual', `http://stca.tas.gov.au/weeds/wp-content/uploads/2010/01/weed-management-manual-.pdf`. [Online; accessed 17th June-2012].

Bai, X. D., Cao, Z. G., Wang, Y., Yu, Z., Zhang, X. & C.N.Li (2013), 'Crop segmentation from images by morphology modeling in the cie l*a*b* color space', *Comput. Electron. Agric.* **99**, 21–34.

Baron, R., Cro, W. & Wolf, T. (2002), Dual camera measurement of crop canopy using reflectance, *in* 'AIC 2002 Meeting', Vol. 02-209, Canadian Society of Agricultural Engineering.

Berge, T., Goldberg, S., Kaspersen, K. & Netland, J. (2012), 'Towards machine vision based site-specific weed management in cereals', *Computers and Electronics in Agriculture* **80**, 79–86.

Bernat, G., Burns, A. & Llamos, A. (2001), 'Weakly hard real-time systems', *IEEE Transactions on computers* **50(4)**, 308–321.

Berry, G. (1989), Real time programming: Special purpose or general purpose languages, *in* 'Proceedings of the IFIP 11th World Computer Congress', pp. 11–17.

Blatt, M., Wiseman, S. & Domany, E. (1996), 'Superparametric clustering of data', *Physical Review Letters* **76(18)**.

Bo, Z., Zongjia, W., Peng, Z., Wenhua, M. & Xiaochao, Z. (2012), 'Design and experiment of intelligent weed recognition systems', *Transactions of*

the Chinese Society of Agricultural Engineering (Transactions of the CSAE) **28(Supp.2)**, 184–187.

Borregaard, T., Nielsen, H., Norgaard, L. & Have, H. (2000), 'Crop-weed discrimination by line imaging spectroscopy', *Journal Agricultural Engineering Research* **75**, 389–400.

Bossu, J., Gee, C., Jones, G. & Truchetet, F. (2009), 'Wavelet transform to discriminate between crop and weed in perspective agronomic images', *Computers and Electronics in Agriculture* **65**, 133–143.

Bracewell, R. (1965), *Convolution and Two-dimensional convolution, in The Fourier Transform and Its Applications (pp. 2550 and 243244)*, Published by McGraw-Hill,NewYork.

Bradski, G. & Kaehler, A. (2008), *Learning OpenCV*, Published by OReilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

Brodatz, P. (1966), *"Textures: A Photographic Album for Artists and Designers"*, Dover Publications, New York.

Burks, T., Shearer, S. & Payne, F. (2000), 'Classification of weed species using colour texture features and discriminant analysis', *ASAE* **43(2)**, 441–448.

Callow, B., Fillols, E. & Willcox, T. (2010), *Weed management manual*, Bureau of Sugar Extension Services (BSES). BSES limited technical publication MN10004.

Campbell, C. & Miller, A. (2010), *Parallel programming with Microsoft visual c++*, Microsoft Corporation. Published free and online `http://msdn.microsoft.com/en-us/library/gg675934.aspx`.

Cerutti, G., Tougnel, L., Mille, J., Vacavant, A. & Coquin, D. (2011), Guiding active contours for tree leaf segmentation and identification, *in* 'Proceedings from the Second International Conference of the Cross-Language Evaluation Forum', Cross-language Evaluation Forum.

Chalak, M., A, R. & Ierland, E. C. (2011), 'Biological control of invasive plant species: A stochastic analysis', *Weed Biology and Management* **11**, 137–151.

Chapron, M., Requena-Esteso, M., Boissard, P. & Assemat, L. (1999), A method for recognizing vegetable species from multispectral images, *in* 'In: Precision Agriculture 1999 2nd European Conference on Precision Agriculture', Vol. Part 1 (ed. JV Stafford), 239-248, Sheffield Academic Press, Sheffield, UK.

Chene, Y., Rousseau, D., Lucidarme, P., Bertheloot, J., Caffier, V., Morel, P., Belin, E. & Chapeau-Blondeau, F. (2012), 'On the use of depth camera for 3D phenotyping of entire plants', *Computers and Electronics in Agriculture* **82**, 122–127.

Chuang, Y., Curless, B., Salesin, D. & Szeliski, R. (2001), A bayesian approach to digital matting, *in* 'Proceedings of IEEE CVPR 2001', Vol. 2, IEEE Computer Society, pp. 264–271.

Copea, J. S., Corneyb, D., Clarkb, J. Y., Remagninoa, P. & Wilkinc, P. (2012), 'Autonomous robotic weed control systems : a review', *Expert Systems with Applications* **39(8)**, 75627573.

Corke, P. (1996), *Visual control of robots : high-performance visual servoing*, Robotics and mechatronics series, Research Studies Press ; Wiley, Taunton England ; New York, NY.

Cotton CRC Australia (2010), 'Cotton pest management guide 2010-2011', `http://www.cottoncrc.org.au/industry/Publications/Cotton_Pest_Management_Guide_2010__11`. [Online; accessed 12th June 2012].

Croplands Australia Ltd (2013), 'Product catalogue', `http://www.croplands.com.au`. [Online; accessed 10th June 2013].

Cropoptics Australia (2012), 'Grower information note', `http://www.cropoptics.com.au`. [Online; accessed 10th June 2012].

da Silva Pires, D., Cesar, R., Vieira, M. & Velho, L. (2005), Tracking and matching connected components from 3D video, *in* 'Computer Graphics and Image Processing, 2005. SIBGRAPI 2005. 18th Brazilian Symposium on', pp. 257–264.

Dart, I. (2013), 'sugarcane density estimation', verbal. Ian Dart, farm manager, Fairymead farm, Bundaberg.

De-Rainville, F.-M., Durand, A., Fortin, F.-A., Tanguy, K., Maldague, X., Panneton, B. & Simard, M.-J. (2014), 'Bayesian classification and unsupervised learning for isolating weeds in row crops', *Pattern Analysis and Applications* **17**(2), 401–414.

Department of Primary Industries NSW (2009), 'Consevation farming', `http://www.western.cma.nsw.gov.au`. [Grower information note].

Dorigo, M. & Blum, C. (2005), 'Ant colony optimization theory, a survey', *Theoretical Computer Science* **355,issues 2-3**, 243–278.

Dorrington, A. (2014), 'email communication', `http://www.chronoptics.com/`. Chronoptics group at the University of Waikato, New Zealand.

El-Faki, M., Zhang, N. & Peterson, D. (2000), 'Factors affecting color-based weed detection.', *Transactions of the ASAE* **43**, 1001–1009.

Emmert, E. (1881), 'Grenverhltnisse der nachbilder. klinische monatsbltter fr augenheilkunde und fr augenrztliche fortbildung', - **19**, 443–450.

Encyclopedia Britannica (2013), 'Fallow farming system', `http://www.britannica.com/EBchecked/topic/200948/fallow-system`. [Online; accessed 10th March 2013].

Fairchild, M. D., Rosen, M. R. & Johnson, G. M. (2014), 'Spectral and metameric color imaging', `http://www.cis.rit.edu/mcsl/research/PDFs/Spec_Met.pdf`. [Online; accessed 29th December 2014].

Fan, J., Yau, D., Elmagarmid, A. & Aref, W. (2001), 'Automatic image segmentation by integrating color-edge extraction and seeded region growing', *Image Processing, IEEE Transactions on* **10**(10), 1454–1466.

Fergal, O. (2010), 'Striking the balance 2nd edition', `http//www.nt.gov.au/d/Primary_Industry/index.cfm?Header=Striking%20the%20Balance%20-%202nd%20Edition`. [chapter 3,Online; accessed 11th November 2012].

FitPc (2013), 'Fit-pc homepage', `http://www.fit-pc.com/web/`. [Online; accessed 29-9-2013].

Franz, E., Gebhardt, M. R. & Unklesbay, K. B. (1991), 'Shape description of completely visible and partially occluded leaves for identifying plants in digital images', *Transactions of the ASAE* **34(2)**, 673–681.

Gee, C., Bossu, J., Jones, G. & Truchetet, F. (2006), 'Crop/weed discrimination in perspective agronomic images', *Computers and electronics in agriculture* **60**, 4959.

Gerhards, R. & Chrisensen, S. (2003), 'Real-time weed detection, decision making and patch spraying in maize, sugarbeet, winter wheat and winter barley', *Weed Research* **43**, 385–392.

Ghazali, K., Mustafa, M. M. & Hussain, A. (2008*a*), Machine vision system for automatic weeding strategy in oil palm plantation using image filtering technique, *in* 'Information and Communication Technologies: From Theory to Applications, 2008. ICTTA 2008. 3rd International Conference on', Advanced technologies, pp. 1–5.

Ghazali, K., Mustafa, M. M. & Hussain, A. (2008*b*), 'Machine vision system for automatic weeding strategy using image processing technique', *American-Eurasian Journal of Agricultue and Environmental Science* **3(3)**, 451–458.

Gliever, C. & Slaughter, D. (2001), Crop versus weed recognition with artificial neural networks, *in* 'Session 95:Real-Time Image Applications, International meeting of ASAE 2001', Vol. Paper 01-3104.

Golzarian, M. & Frick, R. (2011), 'Classification of images of wheat, ryegrass and brome grass species at early growth stages using principal component analysis', *Plant Methods* **7:28**.

Greenspan, H., Belongie, S., Goodman, R. & Perona, P. (1994), Rotation invariant texture recognition using a steerable pyramid, *in* 'Pattern Recognition, 1994. Vol. 2 - Conference B: Computer Vision amp; Image Processing., Proceedings of the 12th IAPR International. Conference on', Vol. 2, pp. 162–167 vol.2.

Grossberg, S. (1987), 'Competitive learning: From interactive activation to adaptive resonance', *Cognitive Science* **11**, 23–63.

Guerrero, J., Pajares, G., Montalvo, M., Romeo, J. & Guijarro, M. (2012), 'Support vector machines for crop/weeds identification in maize fields', *Expert Systems with Applications* **39**, 11149–11155.

Guijarro, M., Pajares, G., Riomoros, I., Herrera, P., Burgos-Artizzu, X. & Ribeiro, A. (2011), 'Automatic segmentation of relevant textures in agricultural images.', *Computers and Electronics in Agriculture* **75**, 75–83.

Guo, Z., Zhang, L. & Zhang, D. (2010), 'Rotation invariant texture classification using lbp variance (lbpv)with global matching', *PatternRecognition* **43**, 706–719.

Haggar, R., Stent, C. & Isaac, S. (1983), 'A prototype hand-held patch sprayer for killing weeds, activated by spectral differences in crop/weed canopies.', *Journal of Agricultural Engineering Research* **28**, 49–358.

Halbwachs, N. (1993), *Synchronous Programming of Reactive Systems*, Kluwer Academic Publishers, Dordrecht, Netherlands. BSES limited technical publication MN10004.

Hall-Beyer, M. (2013), 'The grey level co-occurrence matrix,current version: 2.10 february 2007', `http://www.fp.ucalgary.ca/mhallbey/the_glcm.htm`. [Online; accessed 6th June 2013].

Haralick, R. (1979), 'Statistical and structural approaches to texture', *Proceedings of the IEEE* **67(5)**, 786–804.

Haralick, R. M. & Shapiro, L. G. (1985), Image segmentation techniques, *in* 'Proc. SPIE', pp. 2–9.

Hartigan, J. A. (1975), *Clustering Algorithms (Probability and Mathematical Statistics)*, John Wiley and Sons Inc.

Hartmann, G. (1987), 'Recognition of hierarchically encoded images by technical and biological systems', *Biological Cybernetics* **57**, 73–84.

Haug, S., Michaels, A., Biber, P. & Ostermann, J. (2014), Plant classification system for crop /weed discrimination without segmentation, *in* 'Applications of Computer Vision (WACV), 2014 IEEE Winter Conference on', pp. 1142–1149.

Hemming, J. & Rath, T. (2001), 'Computer-vision-based weed identification under field conditions using controlled lighting', *Journal of Agricultural Engineering Research* **78 (3)**, 233–243.

Hill, M. & Marty, M. (2008), 'Amdahl's law in the multicore era', *Computer (IEEE)* **41**(7), 33–38.

Hong, S., Minzan, L. & Zhang, Q. (2012), 'Detection system of smart sprayers-status, challenges and perspectives', *International Journal of Agricultural and Biological Engineering* **5(3)**, 10–23.

Intel (2003), 'Hyper-threading technology technical users guide', `http://cache-www.intel.com/cd/00/00/01/77/17705_htt_user_guide.pdf`. [Online; accessed 11-3-14].

Intel (2012), 'Multi-core introduction', `http://software.intel.com/en-us/articles/multi-core-introduction`. [Online; accessed 20th March-2013].

Jafari, A., Mohtasebi, S., Jahromi, H. & Omid, M. (2006), 'Weed detection in sugar beet fields using machine vision', *International Journal of Agricultural and Biological Engineering* **8(5)**.

Jeon, H., Tian, L., T.Grift, L.Bode & Hagar, A. (2009), Stereovision system and image processing algorithms for plant specific application, *in* 'ASABE Annual International Meeting', Vol. Paper Number: 090047.

Jeon, H. Y., Tian, L. F. & Zhu, H. (2011), 'Robust crop and weed segmentation under uncontrolled outdoor illumination', *Sensors* **11**, 6270–6283.

Ji, B., Zhu, W., Liu, B., Ma, C. & Li, X. (2009), Review of recent machine-vision technologies in agriculture, *in* 'Second International Symposium on Knowledge Acquisition and Modeling', IEEE Computer Society.

Jin, J. & Tang, L. (2009), 'Corn plant sensing using real-time stereo vision', *Journal of Field Robotics* **26(67)**, 591–608.

Keranen, M., Aro, E., Tyystjarvi, E. & Nevalainen, O. (2003), 'Automatic plant identification with chlorophyll fluorescence fingerprinting', *Precision Agriculture* **4**, 53–67.

Kinect (2013), 'Kinect homepage', `http://www.xbox.com/en-US/KINECT`. [Online; accessed 29-9-2013].

Kinmont, A., Rees, J. & Rees, S. (1999), 'Patent us5924239; controller for agricultural sprayers', `www.google.com/patents/US5924239`. [Online; accessed 17th November-2013].

Laksberg, A. (2012), 'Asynchronous programming in c++ using ppl', *MSDN Magazine* **27(2)**, 1. viewed online 11-3-14 `http://msdn.microsoft.com/en-us/magazine/hh781020.aspx`.

Lamm, R., Slaughter, D. & Giles, D. (2002), 'Precision weed control system for cotton', *Transactions of the ASAE* **45(1)**, 231–238.

Langer, H., Bottger, H. & Schmidt, H. (2006), 'A special vegetation index for the weed detection in sensor based precision agriculture', *Environmental Monitoring and Assessment* **117**, 505–518.

LeCun, Y., Bottou, L., Orr, G. & Muller, K. (1998), *Efficient BackProp, in G. Orr and K. Muller (Eds.), Neural Networks: Tricks of the Trade*, New York: Springer-Verlag.

Li, X. & Chen, Z. (2010), Weed identification based on shape features and ant colony optimization algorithm, *in* 'International Conference on Computer Application and System Modeling (ICCASM 2010)', Vol. V1, IEEE.

Lin, K. & Burke, E. (1992), 'Coming to grips with real-time realities', *IEEE Software* **9(1)**, 12–15.

Longchamps, L., Panneton, B., Samson, G., Leroux, G. & Thriault, R. (2009), 'Discrimination of corn, grasses and dicot weeds by their UV-induced fluorescence spectral signature', *Precision Agriculture* **11(2)**, 181–197.

Lowel Education Centre (2012), 'Color temperature and color rendering index demystified', `www.lowel.com/edu/color_temperature_and_rendering_demystified.html`. [Online; accessed 11th November 2012].

Materka, A. & Strzelecki, M. (1998), Texture analysis methods-a review, *in* 'COST B11 report', Institute of Electronics, Technical University of Lodz, Poland.

McCarthy, C., Rees, S. & Baillie, C. (2012), Preliminary evaluation of shape and colour image sensing for automated weed identification in sugarcane., *in* 'In: 34th Annual Conference of the Australian Society of Sugar Cane Technologists (ASSCT 2012), 1-4 May 2012, Cairns, Australia.', Australian Society of Sugarcane Technologists.

Metz, C. (2006), 'Receiver operating characteristic analysis: A tool for the quantitative evaluation of observer performance and imaging systems', *Journal American College of Radiology* **3**, 413–422.

Meyer, G., Mehta, T., Kocher, M., Mortensen, D. & Samal, A. (1999), 'Textural imaging and discriminant analysis for distinguishing weeds for spot spraying', *Transactions of the ASAE* **41(4)**, 1189–1197.

Mishra, A., Matou, K., Mishra, K. & Nedbal, L. (2009), 'Towards discrimination of plant species by machine vision-advanced statistical analysis of chlorophyll fluorescencetransients', *Journal Fluorescence* **19**, 905–913.

MSDN, M. (2015), 'Setthreadpriority function', `https://msdn.microsoft.com/en-us/library/windows/desktop/ms686277(v=vs.85).aspx`. Online; accessed 2-6-2015.

Nakarmia, A. & Tang, L. (2012), 'Automatic inter-plant spacing sensing at early growth stages using a 3D vision sensor', *Computers and Electronics in Agriculture* **82**, 23–31.

Nieuwenhuizen, A., Hofstee, J., van de Zande, J., Meuleman, J. & van Henten, E. (2010), 'Classification of sugar beet and volunteer potato reflection spectra using a neural network to select discriminative wavelengths', *Transactions of the ASAE* **73(2)**, 146–153.

Nilsson, N. J. (1998), *MACHINE LEARNING*, Stanford University, USA (online). `http://robotics.stanford.edu/people/nilsson/mlbook.html`.

Noble, S. & Brown, R. (2002), The use of spectral properties for weed detection and identification   a review, *in* 'CSAE Paper number 02-208', Canadian Society of Agricultural Engineering.

Noble, S. & Crowe, T. (2001), Plant discrimination based on leaf reflectance, *in* 'Annual International Meeting of ASAE', Vol. Paper No. 01-1150, ASAE.

Ogale, A. & Aloimonos, Y. (2005), 'Shape and the stereo correspondence problem', *International Journal of Computer Vision* **65 (1)**.

Ojala, T., Pietikainen, M. & Maenpaa, T. (2002), 'Multiresolution gray-scale and

rotation invariant texture classification with local binary patterns', *IEEE Transactions on pattern analysis and machine intelligence* **24 (7)**, 971–987.

Okamoto, H., Murata, T., Kataoka, T. & Hata, S. (2007), 'Plant classification for weed detection using hyperspectral imaging with wavelet analysis.', *Weed Biology and Management* **7**, 31–37.

OpenCV Devzone (2013), 'Opencv overview', `http://code.opencv.org/projects/opencv/wiki`. [Online; accessed 23rd May-2013].

OpenKinect (2013), 'Open kinect overview', `http://openkinect.org/wiki/Main_Page`. [Online; accessed 23rd May-2013].

Otsu, N. (1979), 'A threshold selection method from gray-level histograms', *IEEE Transactions on Systems, Man, and Cybernetics* **9 (1)**, 62–66.

Paproki, A., Fripp, J., Salvado, O., Sirault, X., Berry, S. & Furbank, R. (2011), Automated 3D segmentation and analysis of cotton plants, *in* 'Digital Image Computing Techniques and Applications (DICTA), 2011 International Conference on', pp. 555–560.

Payero, J., Neale, C. & Wright, J. (2004), 'Comparison of eleven vegetation indexes for estimating plant height of alfalfa and grass.', *American Society of Agricultural Engineers* **20(3)**, 385–393.

Peng, Z. (2011), 'Robust weed recognition using blur moment invariants', *Biosystem Engineering* **110**, 198–205.

Perez, A., Lopez, F., Benlloch, J. & Christensen, S. (2000), 'Colour and shape analysis techniques for weed detection in cereal fields', *Computers and Electronics in Agriculture* **25**, 197–212.

Persa, S., Nicolescu, C. & Jonker, P. (2000), Evaluation of two real time low level image processing architectures, *in* 'Proceedings from the IAPR Workshop on Machine Vision Applications', pp. 295–298.

Persson, M. & Astrand, B. (2008), 'Classification of crops and weeds extracted by active shape models', *Biosystems Engineering* **100**, 484–497.

Piron, A., Leemans, V., Kleynen, O., Lebeau, F. & Destain, M. (2008), 'Selection of the most efficient wavelength bands for discriminating weeds from crop', *Computers and Electronics in Agriculture* **62**, 141–148.

Piron, A., Leemans, V., Lebeau, F. & Destain, M. (2009), 'Improving in-row weed detection in multispectral stereoscopic images', *Computers and Electronics in Agriculture* **69**, 73–79.

Piron, A., van der Heijden, F. & Destain, M. F. (2011), 'Weed detection in 3D images', *Precision Agriculture* **12**, 607–622.

Plotze, R., Falvo, M., Padua, J. G., Bernacci, L. C., Vieira, M. L. C., Oliveira, G. C. X. & Martinez, O. (2005), 'Leaf shape analysis using the multiscale minkowski fractal dimension, a new morphometric method: A study with passiora (passi oraceae)', *Canadian Journal Of Botany* **83**, 287–301.

Primary Industries Standing Committee 82 (2002), 'Spray drift management, principles, strategies and supporting information', CSIRO publishing, 150 Oxford street, Collingwood Victoria, Australia.

Prime Sense (2013), 'Prime sense website', `http://www.primesense.com/`. [Online; accessed 23rd May-2013].

Queensland Department of Agriculture, F. & Forests (2012), 'Ways to prevent weed spread', `http://www.daff.qld.gov.au/4790_7070.htm`. [Online; accessed 19th November-2012].

Queensland Department of Primary Industries & Fisheries (2005), 'Agriculture chemical user's manual', `http://www.daff.qld.gov.au/documents/Biosecurity_AgVetChemicalsAndResidues/AgChem-UsersManual.pdf`. [Online; accessed 2nd December-2012].

Rees, S., Burgos-Artizzu, X., McCarthy, C. & Baillie, C. (2009), 'Development of a prototype precision spot spray system using image analysis and plant identification technology', *National Centre for Engineering in Agriculture Publication 1002306/1, USQ, Toowoomba* .

Reid, R. (1990), *A Manual of Australian agriculture*, Australian Institute of Agricultural Science.

Relf, C. (2004), *Image Acquisition and Processing with LabView.*, CRC Press, Abingdon, U.K.

Romeo, J., Pajares, G., Montalvo, M., Guerrero, J., Guijarro, M. & de la Cruz, J. (2013), 'A new expert system for greenness identification in agricultural images', *Expert Systems with Applications* **40**, 22752286.

Rong, L. (July 2010), A run-based algorithm for identifying connected components in binary images, *in* 'Software Engineering and Service Sciences (IC-SESS), 2010 IEEE International Conference on', pp. 137–141.

Sabeenian, R. & Palanisamy, V. (2009), 'Texture based weed detection using multi resolution combined statistical and spatial frequency (MRCSF)', *World Academy of Science, Engineering and Technology* **52**.

Sachez, A. & Marchant, J. (2000), Fusing 3D information for crop/weeds classification, *in* '15th International Conference on Pattern Recognition (ICPR'00)', Vol. 4, pp. 42–95.

Sansoni, G., Trebeschi, M. & Docchio, F. (2009), 'State-of-the-art and applications of 3D imaging sensors in industry, cultural heritage, medicine, and criminal investigation', *Sensors* **9**, 568–601.

Scott, J. & Farqhuarson, R. (2004), 'An assessment of the economic impacts of NSW agriculture research and extension: Conservation farming and reduced tillage in northern nsw', *Economic Research Report* **19**.

Seatovic, D., Grninger, R., Thomas, A. & Martin, H. (2008), 3D object recognition, localization and treatment of rumex obtusifolius in its natural environment, *in* '1st International Conference on Machine Control and Guidance'.

Shroeder, B. (2014), 'Sugarcane weed destruction', verbal. Bernard Shroeder is the NCEA sugarcane specialist.

Slaughter, D., Giles, D. & Downey, D. (2008), 'Autonomous robotic weed control systems : a review', *Computers and Electronics in Agriculture* **61**, 63–78.

Software, S. (2014), 'The Trigger Modes For Area Scan Cameras', `http://www.silicon-software.info/phocadownload/download/live_docu/RT5/en/documents/TriggerSystem/TriggerModesArea.html`. [Online; accessed 22-12-2014].

Sogaard, H. (2005), 'Weed classification by active shape models', *Biosystems Engineering* **91(3)**, 271–281.

Souza, J., Batista, R., Moreira, D. & Leite, C. (2008), Identification of weeds based in the texture characteristics, *in* 'International Conference of Agricultural Engineering,', CIGR, Conbea, Brazil, August 31 to September 4, 2008.

Stankovic, J. (1992), 'Real time computing', *Byte* **17 (8)**, 155–162.

Steward, B. & Tian, L. (1998), 'Real-time weed detection in outdoor field conditions', *Precision Agriculture and Biological Quality* pp. 266–278.

Suh, H. K., Hofstee, J. W. & van Henten, E. J. (2014), Shadow-resistant segmentation based on illumination in-variant image transformation, *in* 'International Conference of Agricultural Engineering', Vol. Ref C0475, Proceedings International Conference of Agricultural Engineering, Zurich, 06-10.07.2014.

Swain, K., Nrremark, M., Jrgensen, R., Midtiby, H. & Green, O. (2011), 'Weed identification using an automated active shape matching (aasm) technique', *Biosystems Engineering* **110(4)**, 450–457.

Tang, L., Tian, L. & Steward, B. (2000), 'Color image segmentation with genetic algorithm for in-field weed sensing.', *Transactions of the ASAE* **43**, 1019–1027.

Tang, X. (1998), 'Texture information in run-length matrices', *IEEE Transactions on image processing* **7(11)**, 1602–1609.

Tangwongkit, R., Salokhe, V. & Jayasuriya, H. (2006), 'Development of a tractor mounted real-time,variable rate herbicide applicator for sugarcane planting', *Agricultural Engineering International Ejournal* **Vol.8**, Manuscript PM 06 009.

Targhi, T., Marten, B., Eric, H. & Jan-Olof, E. (2006), Real-time texture detection using the lu-transform, *in* 'Computation Intensive Methods for Computer Vision Workshop', ECCV06 , Graz, Austria, May 2006.

The United Nations Food & Agricultural Organisation (2012), 'Conservation agriculture', `http://www.fao.org/ag/ca/`. [Online; accessed 12th June-2012].

Tian, L., Reid, J. & Hummel, J. (1999), 'Development of a precision sprayer for site specific weed management', *Transactions of the ASAE* **42(4)**, 893–900.

Tian, L., Slaughter, D. & Norris, R. (1997), 'Outdoor field machine vision identification of tomato seedlings for automated weed control', *Transactions of the ASAE* **40(6)**, 1761–1768.

Tuceryan, M. & Jain, A. (1998), *The Handbook of Pattern Recognition and Computer Vision (2nd Edition)*, World Scientific Publishing Co.

University of Minnesota (2008), 'Herbicide resistance', `http://www.extension.umn.edu/distribution/cropsystems/components/DC6077a.html`. [Online; accessed 20th March-2013].

Vapnik, V. (1995), *The Nature of Statistical Learning Theory*, New York: Springer-Verlag.

Wallenberg, M., Felsberg, M., Forssen, P.-E. & Dellen, B. (2011), *Channel Coding for Joint Colour and Depth Segmentation*, Vol. 6835, Springer Berlin Heidelberg.

Wang, N., N.Zhang, Dowell, F., Sun, Y. & Peterson, D. (2001), 'Design of an optical weed sensor using plant spectral characteristics', *Transactions of the ASAE* **44(2)**, 409–419.

Wang, N., N.Zhang, J.Wei, Stoll, Q. & Peterson, D. (2007), 'A real-time, embedded, weed-detection system for use in wheat fields', *Biosystems Engineering* **98(3)**, 276–285.

Weiss, U. & Biber, P. (2011), 'Plant detection and mapping for agricultural robots using a 3D lidar sensor', *Robotics and Autonomous Systems* **59**, 265–273.

Weiss, U., Biber, P., Laible, S., Bohlmann, K. & Zell, A. (2010), Plant species classification using a 3D lidar sensor and machine learning, *in* 'Machine Learning and Applications (ICMLA), 2010 Ninth International Conference on', pp. 339–345.

West, P. C. (2001), 'High speed,real-time machine vision', online`http://www.imagenation.com/pdf/highspeed.pdf`. [Online; accessed jan-2015].

Wilson, C. (2013), 'Dust bowl', `https://prezi.com/bzkusaqpsw8f/dust-bowl/`. [Online; accessed 23rd May-2013].

Woebbecke, D., Meyer, G., Bargen, K. V. & Mortensen, D. (1995), 'Color indeces for weed identification under various soil, residue, and lighting conditions', *Transactions of the ASAE* **38(1)**, 259–269.

Wu, L., Wen, Y., Deng, X. & Peng, H. (2009), 'Identification of weed-corn using bp network based on wavelet features and fractal dimension', *Scientific Research and Essay* **4 (11)**, 1194–1200. [http://www.academicjournals.org/SRE; accessed jan-2015].

Wu, X., W.Xu, Song, Y. & Cai, M. (2011), 'A detection method of weed in wheat field on machine vision', *Procedia Engineering* **15**, 1998 2003.

Wylie sprayers Inc (2013), 'Product catalogue', `http://www.wyliesprayers.com`. [Online; accessed 10th June 2013].

Zhang, D. (2004), The optimality of naive bayes, *in* 'Proceedings of the 17th International The Florida Artificial Intelligence Research Society (FLAIRS 2004) Conference', Miami Beach, Florida, May 17-19, 2004.

Zhang, H., Fritts, J. & Goldman, S. (2008), 'Image segmentation evaluation: A survey of unsupervised methods', *Computer Vision and Image Understanding* **110**, 260–280.

Zhang, Y., Slaughter, D. & Staab, E. (2012), 'Robust hyperspectral vision-based classification for multi-season weed mapping', *ISPRS Journal of Photogrammetry and Remote Sensing* **69**, 65–73.

Appendix A

Glossary of terms

Term	Definition
BST	Binarised segmentation technique ($G > R$ and $G > B$)
DATE FORMAT	day/month/year
DCS	Depth, colour and size (algorithm)
DCSA	Depth colour segmentation (algorithm)
DCSS	Depth, colour, size and spatial (algorithm)
Fallow	Land that is left unseeded with crop and weed free during a growing season .
GLCM	Grey level co-occurrence matrix (texture features)
GLRLM	Grey level run length matrix (texture features)
k-NN	K nearest neighbour (classifier)
LBP	Local binary pattern (texture features)
LBPD	Local binary patter and depth (custom texture features)
Min-till	Minimum tillage farming system
MLP	Multi layer perceptron (neural network classifier)
NIR	Near infra-red portion of the spectrum.
No-till	No tilling farming system
Occlusion	An occlusion is where parts of plants and/or stubble overlap each other from the viewpoint of the camera.

OIR	Object identification redundancy
OTC	Object tracking classification
Ratoon	Crop regrowth season
Real-time	A computation result must be correct and within a strict deadline
Real-world	Actual commercial field conditions
ROC	Receiver operating characteristic curves (feature selection tool)
ROI	Region within the image containing the area of interest for machine vision.
SVM	Support vector machine (classifier)
SP	Spatial position (algorithm)
Spot spray	The spot application of a herbicide to an individual or small patch of weed or crop.
SPP	Synchronised parallel processing technique
Trash blanket	Sugarcane weed suppression practice
VIS	Visible portion of the spectrum.
Weed mapping	The recording of the weed and its latitude and longitude so that it can be displayed on geographic information software for evaluation.

Table A.1: Glossary of terms.

Appendix B

Commercial spot sprayers

Weed detection technologies commercially available at the time of writing are only able to distinguish plant from background (i.e. soil and /or stubble). Commercially available technologies are the WeedSeeker® from Trimble corporation and the WeedIT from GPS-Ag Pty Ltd. The Greenseeker® is a similar technology to the WeedSeeker® by Trimble corporation that can be used to determine an 'index' value of the plants beneath the sensor, for nutrient application.

B.1 WeedSeeker® / Greenseeker®

The WeedSeeker® technology (Figure B.1) was released commercially in the 1980s. The WeedSeeker® determines a vegetation index similar to NDVI by measuring the red reflectance and NIR reflectance. These measurements achieve discrimination between green plants and background (i.e. soil) because a green plant has a much higher NIR reflectance than red reflectance, whereas background soil has low reflectance in both NIR and red wavelengths. The WeedSeeker® is an active sensor system providing its own LED light source source to provide illumination for the reflectance reading as well as offsetting ambient light measurement

levels to improve the robustness of the result. The Greenseeker® is an adaptation
of the WeedSeeker® technology which provides a varying vegetation index that
can be used in fertiliser application and/or log measurements to a map. Figure B.1
provides a description of the WeedSeeker® operation and more information can
be accessed at `https://www.trimble.com/Agriculture/WeedSeeker.aspx`.
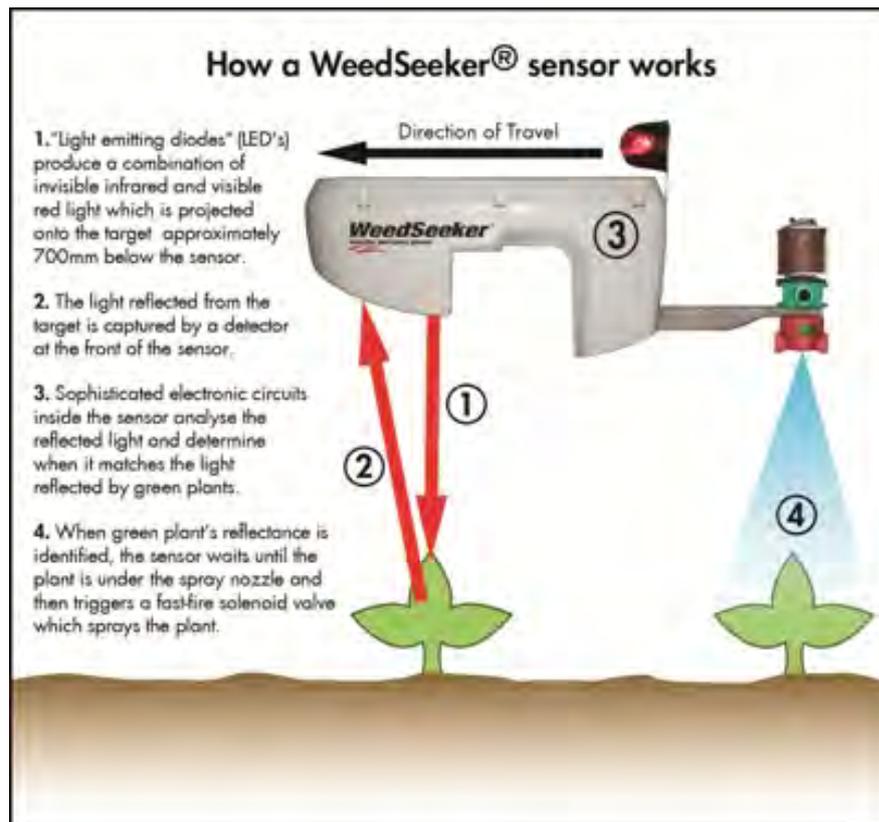
Figure B.1: Diagram of how a WeedSeeker Sensor operates reproduced from
`https://www.trimble.com/Agriculture/WeedSeeker.aspx`.

B.2   WeedIT

WeedIT is a technology that was released in Australia in 2009 by GPS-Ag Pty Ltd and measures NIR emission in response to applied red illumination (Figure 1.9). According to the WeedIT information (`http://www.weedit.com.au/products.php`), the chlorofil in green plants absorbs red light, converts it to NIR and emits it. The WeedIT system has an active red light source and is constantly checking for the emission of NIR above a threshold as those areas have plant material whereas soil and stubble do not have this characteristic.
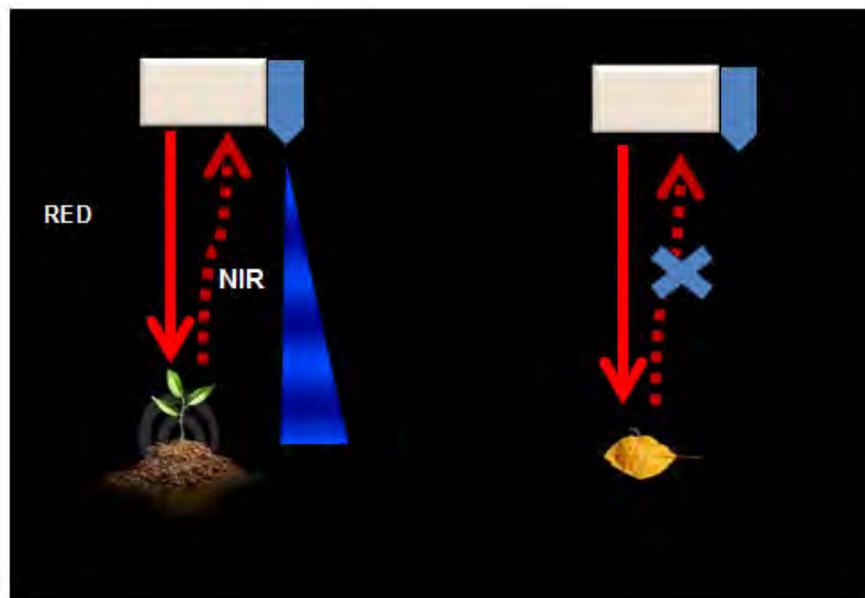
Figure B.2: Diagram of WeedIT sensor operation (reproduced from the WeedIT website `http://www.weedit.com.au`.

B.3 Discussion of commercial weed detection technologies

Commercially-available technologies for weed detection are typically based on point measurements of reflectance/emission of specific wavelengths and make use of the unique properties of plants in the NIR region of the spectrum. Incorporation of a localised and focused artificial light source improves the detection performance in day and night operation but anecdotal evidence from operators (Cropoptics Australia 2012) indicates that the detection systems work most reliably at night in the absence of variable natural daylight. The technologies are most suited to spraying in pre-emergent, inter-row or fallow situations as the detection technique targets any vegetation on a soil or stubble background.

Appendix C

Introduction to sugarcane and pyrethrum farming practices

C.1 Sugar cane farming practices

Sugarcane (*Saccharum officinarum*) in Australia is grown by planting parts of the mature cane plant into rows between 1.4 m to 2 m wide. Sugarcane typically has a growing cycle of three to four 'ratoons' (regrowth crops), growing to a height of between two to four metres. Harvesting is done annually and comprises cutting the cane of at ground level, and chopping it into small (approximately 0.3 m) lengths called billets. The billets are then deposited into haul-out bins to start its journey to the mill where it is crushed (Reid 1990). In minimum and no-till systems, the trash left over from the harvesting process is spread over the ground creating a 'trash blanket' (Figure C.1) up to 0.4 m thick, depending on the amount of vegetative growth in the crop at harvest time. The trash blanket suppresses weed seed germination and is an effective strategy to aid weed control for the industry (Callow et al. 2010).

Grasses like guinea grass and green panic can choke out the sugar cane crop and

Figure C.1: Ratooning sugarcane with a trash blanket.

create non productive areas within the field. Competition between weed and crop (Figure 1.11) can cause significant loss in yield and can shorten the cropping cycle (number of ratoons), which adds to the growing costs, as all farming applications involved in land preparation and planting are amortised over a shorter period, as well as a typically lower yield associated with the first season's harvest.

C.2   Pyrethrum farming practices

Pyrethrum (*Saccharum officinarum*) in Australia is grown from planted seed, as opposed to seedlings and has a perennial crop cycle of up to four years depending on the growing conditions it experiences. The pyrethrum crop is planted in 0.25 m to 0.4 m row spacing with a plant density that causes the row to appear contiguous. Harvesting is achieved by windrowing the pyrethrum and picking it up with modified combine harvesters. The output of the combine harvester contains a significant amount of residue, mixed with the seed, and all is transported to the Botanical Resources Australia facility in Ulverstone, Tasma-

nia, where the residue is removed, and used as a fuel source in the extraction plant. The pyrethrum seeds, which are extremely small, are processed by a proprietary processing plant to extract the pyrethrum oil. All other residues that are extracted in the processing procedure are spread onto the fields as a ground conditioning agent.

After the first year harvest (typically low yielding), farming practices follow a no-till approach, due to the harvest residue left in the fields and the problems associated with tilling the weeds in moist ground, during wet summer and autumn seasons. No-till weed control becomes problematic because of a limited number of herbicides available to control weeds without affecting the pyrethrum. The herbicide limitation is causing weed species to become resistant due to overuse, and some weeds species escape completely, as they are not susceptible to these specific herbicides.

As in sugarcane cropping, the direct competition from weeds can cause reduced yields, as the weeds can choke out the crop reducing the plant density for following seasons. This causes earlier than expected replants, which reduces the amortisation time for the land preparation and planting activities, effectively increasing costs. Precision spray technology would enable the industry to introduce different herbicides into its weed control program, help break resistance, and eradicate the escapes. Figure C.2 shows a tractor and prototype sprayer in first year growing pyrethrum (i.e. not regrowth) and Figure C.3 shows pyrethrum post harvest.

Figure C.2: Tractor and experimental NCEA prototype spot spray system in a crop of first year growing pyrethrum April 2014.

Figure C.3: Data gathering prototype in a crop of post harvest pyrethrum, March 2012.

Appendix D

Patent attorney opinion

University of Southern Queensland   16 April 2013

Summary

In summary, several prior art documents were provided relating to automatic spraying of weeds using image analysis, and further similar prior art documents were identified. The problem of automatic weed spraying thus appears to be a relatively well understood problem in the art.

The use of colour to differentiate weeds and soil, and weeds and crops, appears to be well known in the art. Similarly, the use of depth information in classification of weeds and plant material appears to form part of the prior art base. Thus, such features are not patentable per se.

Additionally, the use of specific line detection algorithms on plant matter appears to be known in the art. While it may be possible to remove one disclosure in this regard from the prior art base in certain jurisdictions, it does not appear that the USQ weed spray system requires use of this specific line detection algorithm.

However the use of depth information together with colour analysis, and in particular the use of depth information together with colour and line based segmentation for the identification of weeds, appears to be novel and inventive in light of the prior art. This appears to have several advantages over the prior art, including a single system that is able to distinguish low grass, sugarcane and foreign objects such as stubble.

It appears that features of the USQ weed spraying system relating to real time processing would be particularly suited as dependent claims in a patent application. These features appear to function independently of the aboveidentified inventive features of the system, and thus do not appear appropriate to be included in broad aspects of the invention.

Accordingly, if the above features are deemed to be of commercial value, we recommend seeking patent protection for your weed sprayer system, and in particular your system including combined depth and colour/line based segmentation.

Please find enclosed for your kind attention our tax invoice in this matter. Yours

sincerely
FISHER ADAMS KELLY
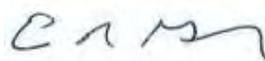
CLINTON PRIDDLE   ERNEST GRAF
Attorney   Partner
MSc   BS JD MIP

Encl: Tax Invoice
        Annexure A

Figure D.1: Summary page of a letter from Fisher, Adams Kelley, patent attorney. The letter reports that the developed segmentation technique and real time process have been found to be inventive and patent protection should be pursued.

Appendix E

Hardware system patent

US005924239A

United States Patent [19]

Rees et al.

[11] **Patent Number:** **5,924,239**

[45] **Date of Patent:** **Jul. 20, 1999**

[54] **CONTROLLER FOR AGRICULTURAL SPRAYERS**

[75] Inventors: **Steven James Rees; James Ian Rees,** both of Croppa Creek; **Andrew Muir Kinmont,** Ferny Hills, all of Australia

[73] Assignee: **Rees Equipment Pty Ltd.,** Australia

[21] Appl. No.: **08/817,841**

[22] PCT Filed: **Oct. 25, 1995**

[86] PCT No.: **PCT/CA95/00595**

§ 371 Date: **Apr. 24, 1997**

§ 102(e) Date: **Apr. 24, 1997**

[87] PCT Pub. No.: **WO96/12401**

PCT Pub. Date: **May 2, 1996**

[30] **Foreign Application Priority Data**

Oct. 25, 1994	[AU]	Australia	....	PM9019
Mar. 17, 1995	[AU]	Australia	....	14946/95
Apr. 26, 1995	[AU]	Australia	....	PN2620
Jun. 23, 1995	[AU]	Australia	....	PN3773

[51] **Int. Cl.^{6}** **A01C 15/00; A01G 15/00**
[52] **U.S. Cl.** **47/1.7**
[58] **Field of Search** 47/1.7; 118/708, 118/707; 250/226, 222.1; 364/550; 404/84.1

[56] **References Cited**

U.S. PATENT DOCUMENTS

3,742,901	7/1973	Johnston	....	118/8
4,144,837	3/1979	Johnston	....	118/8
4,958,306	9/1990	Powell et al.	....	364/550
5,144,767	9/1992	McCloy et al.	....	47/1.7
5,279,068	1/1994	Rees et al.	....	47/1.7

5,294,210	3/1994	Lemelson	....	404/84.1
5,296,702	3/1994	Beck et al.	....	250/226
5,507,115	4/1996	Nelson	....	47/1.7
5,585,626	12/1996	Beck et al.	....	250/222.1
5,606,821	3/1997	Sadjadi et al.	....	47/1.7

FOREIGN PATENT DOCUMENTS

| 41 32 637 | 4/1993 | Germany . |
| 4132637A1 | 4/1993 | Germany . |

OTHER PUBLICATIONS

Saunders,Frederick A., "A Survey of Physics for College Students", Range of Electromagnetic Waves, Chapter 35, 550–566, Sep. 1930.

Primary Examiner—Thomas Price
Assistant Examiner—Fredrick T. French, III
Attorney, Agent, or Firm—Merchant, Gould Smith, Edell, Welter & Schmidt, P.A.

[57] **ABSTRACT**

A controller for agricultural sprayers utilises a detector (**23, 123, 148**) to generate red, blue and green colour signals across a field of view. The colour signals are used to generate a 'its green' or 'not green' output to switch a spray nozzle (**13**) detection of something deemed to be green. The algorithm which determines if there is something which is 'green', rather than 'not green', looks at the level of green component over the red and blue components in the colour signal and if both are exceeded then the decision is that it is 'green'. The level of green over each of red and blue can be compared against preset values to determine the 'green', 'not green' output. The level of green can be established by summing pixel by pixel over an area within the field of view under consideration to see if the sum for the area exceeds a set level to decide that the area is 'green' and requires spraying.

9 Claims, 9 Drawing Sheets
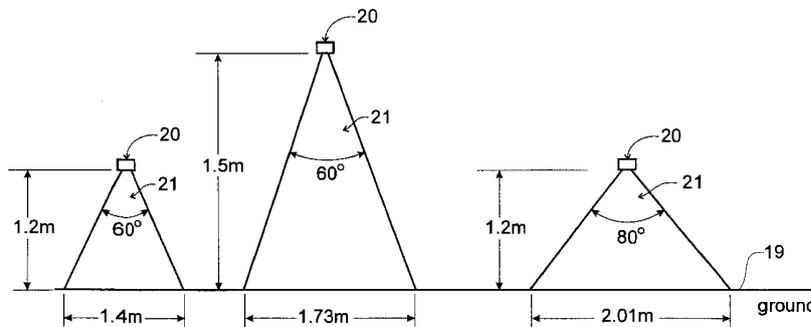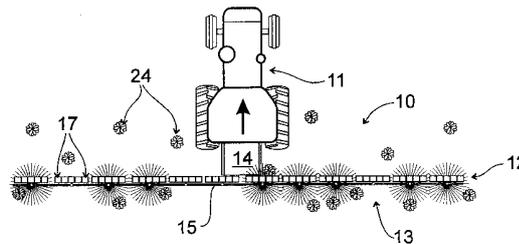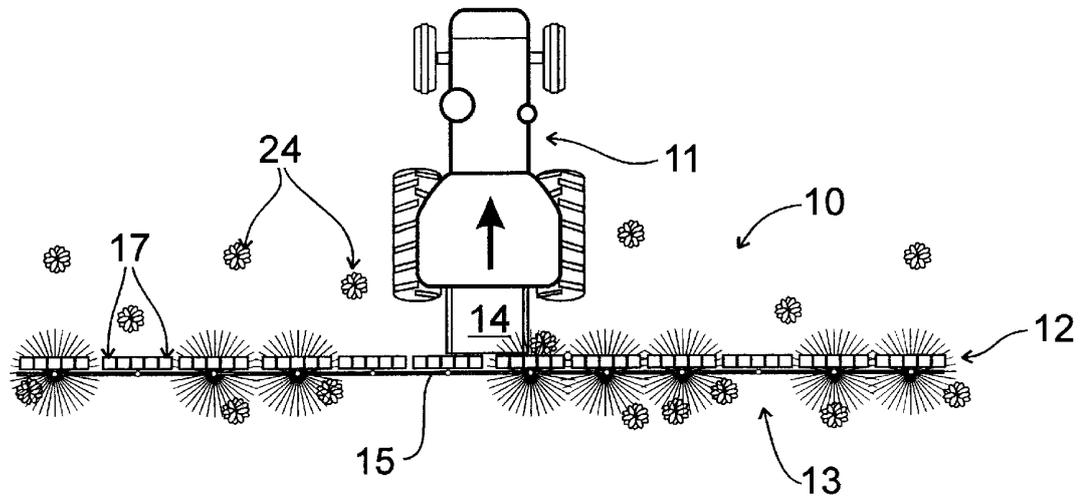
FIG. 1

FIG. 3

FIG. 2C

FIG. 2B

FIG. 2A

CCD — 123

RGB DECODER — 140

RED DIGITISER — 141
GREEN DIGITISER — 142
BLUE DIGITISER — 143

FRAME STORE — 144

R:G:B PROCESSOR — 145

GREEN DISCRIMINATOR — 146

SOLENOID DRIVERS — 147

FIG. 4

FIG. 5

MAIN PROCESS

START ─150

PROCESS SCAN LINE ─151

152─ REGION 1 GREEN?

153─ PREVIOUS SCAN LINE GREEN?

154─ INCREMENT GREEN SCAN LINE COUNTER

155─ COUNT> THRESHOLD ?

156─ SET SOLENOID ON FLAG

157─ CLEAR GREEN SCAN LINE COUNTER

168─ REGION 2 GREEN?

AS ABOVE

169─ REGION 3 GREEN?

AS ABOVE

FIG. 6A

AS ABOVE

170

REGION 4
GREEN?

Y

N

161

FULL
FRAME
?

Y

N

162

SETUP TURN ON AND
TURN OFF TIMES FOR
ALL SOLENOIDS WITH
FLAGS SET

163

SOLENOID
CONTROL
PROCESS

FIG. 6B

151

START — 164

PROCESS REGION — 165

LAST REGION? — 166

N

Y

EXIT — 167

FIG. 7

START — 165

165

WAIT FOR PIXEL — 169

G>R and G>B — 170

N

END REGION? — 172

N

INCREMENT GREEN PIXEL COUNTER — 171

GREEN PIXEL COUNT THRESHOLD — 173

Y

SET FLAG FOR "GREEN" REGION — 174

Y

N

EXIT — 175

FIG. 8

SELENOID CONTROL PROCESS

FIG. 9A

FIG. 9B

CONTROLLER FOR AGRICULTURAL SPRAYERS

FIELD OF THE INVENTION

THIS INVENTION relates to agricultural sprays used to spot spray weeds and the like. In particular the invention relates to a controller by which the spot sprays are selectively activated on determination of the existence of a weed.

BACKGROUND ART

AU-B-37775/89 (618377), the Australian national phase of PCT/AU-89/00267 (WO-89/12510), The Minister for Agricultural and Rural Affairs of the State of New South Wales, discloses a controller for agricultural sprayers where sensors measure the irradiance and radiance (or irradiance and reflectance) of a target area in two bands (eg. red and near infra-red) of the electromagnetic spectrum. The measurements are used to control the spray. Control involves a determination of the relationship between the ratios of the radiance (or reflectance) to the irradiance in each band respectively. The major flaw in this system is that it does not cope with changing light conditions or partly shaded areas in the viewing area. Further it does not provide a size selection function. The plant or weed size at which the controller acts is not able to be adjusted.

Colour analysis is the basis of a variety of discrimination systems operating in a range of circumstances. Examples are seen in U.S. Pat. No. 4,653,014 (Omron) and U.S. Pat. No. 4,797,738 (Tohken). These operate with video signals, operating on components therein to establish the existence of a target condition. In Omron there is seen a totally digital system which uses the R/S, G/S, and B/S signals (where S=R+G+B and R, G, and B are the red, green and blue components of the video signal). This system defines specific colour by analyzing its three signals with reference to upper and lower limits. In Tohken the signals Y (luminance), R–Y and B–Y are compared each with two limit values and analysis determines specific colour. Neither of these systems enables use with sprays in the field where an area which is predominantly green, a weed or other target plant, is to be found in an area of another colour, usually colours such as brown which return a green component in a camera output.

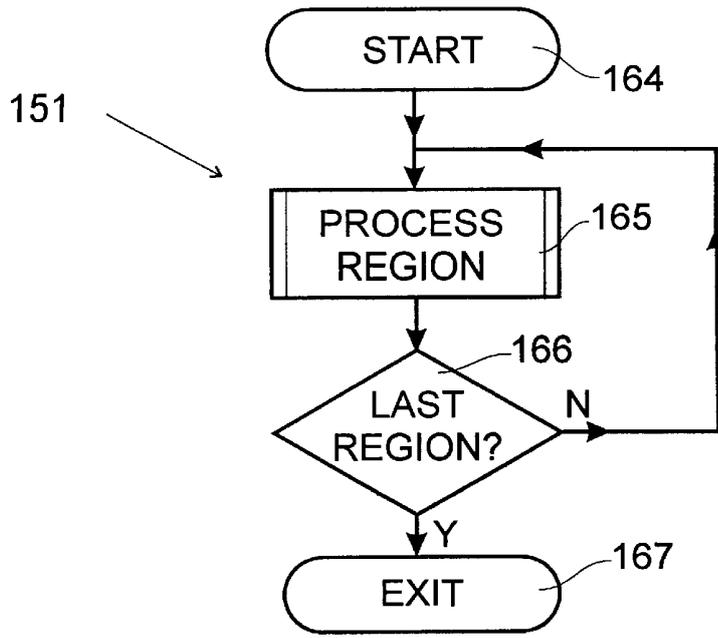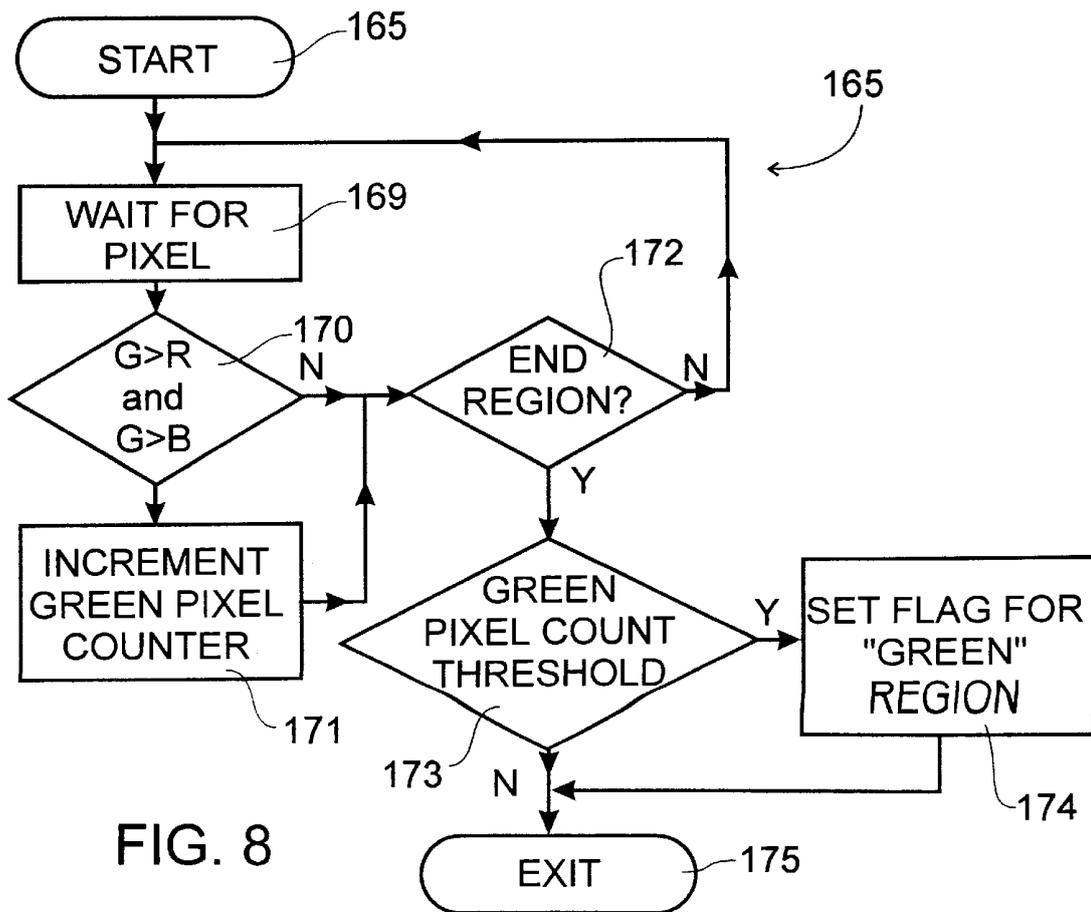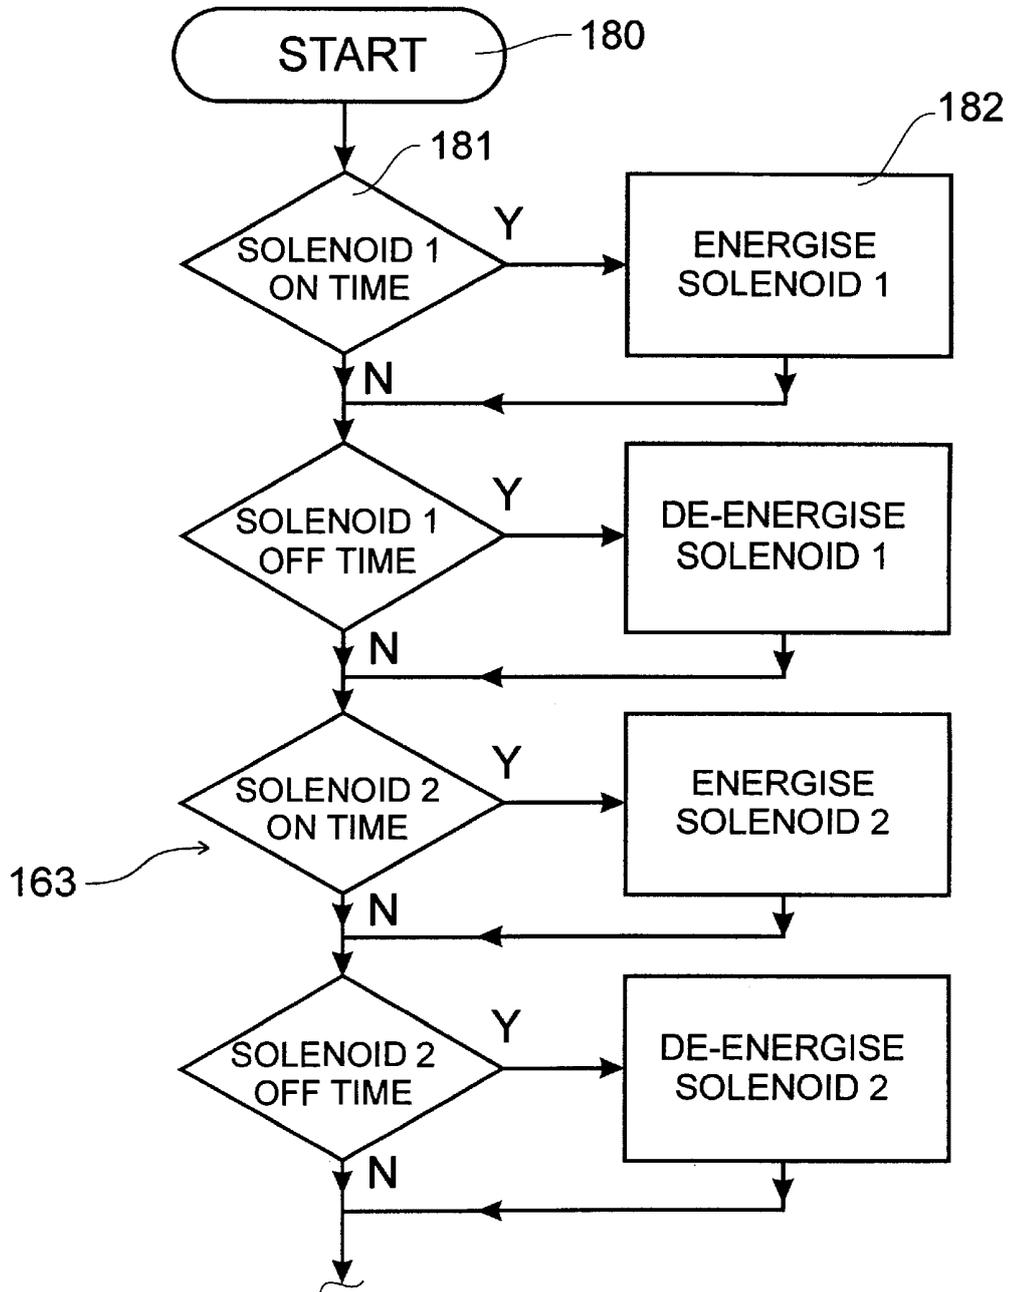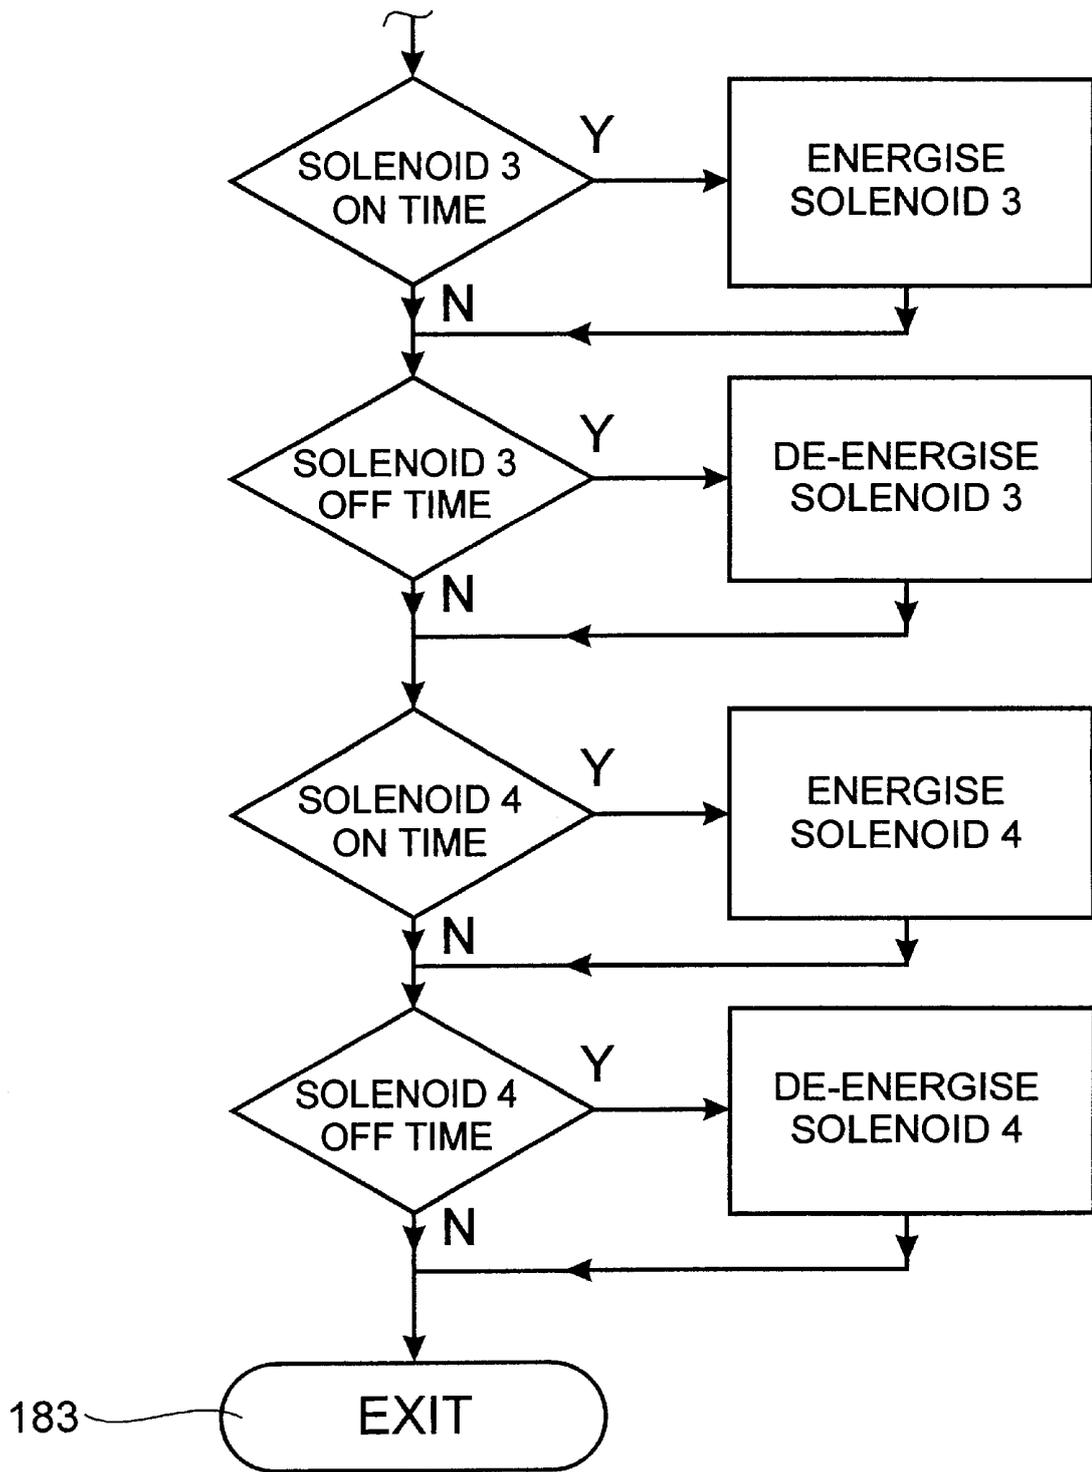## OBJECT OF THE INVENTION

It is an object of the present invention to provide a controller for agricultural sprays, which controller is able to function at normal operational speeds and under varying light conditions, to efficiently locate weeds and other target plants in the field. Other objects and advantages will hereinafter become apparent.

BRIEF SUMMARY OF THE INVENTION

In one aspect, the present invention resides in an agricultural spray controller by which detection of plants on a surface being treated is effected so as to enable the spot application thereto of a spray, said, spray controller comprising:

a spray activation means whereby to action a spray device to effect the spraying of a plant;

a control means for delivering a signal to the spray activation means to effect spraying on detection of a plant;

a detector generating a colour video signal provided in the control means for viewing an area of the surface to be treated and generating an output representative of the

field of view; and control circuitry in the control means coupled to the output of the detector, said control circuitry analyzing the detector output and generating said control signal depending on the detection of a plant;

the control circuitry determining the existence of a plant by examining the colour components of the video signal, noting pixels which are predominantly green, and generating the control signal when the number of predominantly green pixels in an area of the field of view indicates the existence of a green plant.

Evaluation of various plants of interest and their typical backgrounds (soil, rock, stubble, etc) has shown that green foliage has a Green content higher than the Red and Blue content. The same also holds true for the so called colour difference signals, typically denoted as R–Y, B–Y, and G–Y, where Y is luminance. There are some advantages to working with the colour difference signals. The first is that by using the difference signals the effects of ambient light levels can be largely ignored. A second advantage is that CCD cameras with colour difference outputs are more likely to be available. In the ensuing discussion where the system is described without specific reference to luminance either form of signal can be worked with and the alternate form will be readily implemented by the person skilled in the art, there being no special skill required to make the adaptation required to enable use of one rather than the other.

The existence of the green colour of a target weed in the output RGB colour signal of a camera might be determined by a number of processes.

In one form of the invention the Green component of the RGB signal is compared separately to both of the Red and Blue components and if it exceeds both then an 'its green' decision can be made. In a preferred form of this type of controller a suitable selectable offset (setting the level by which the level of green is to exceed the level of red and/or blue) can be introduced so as to allow for different degrees of green of the weeds being treated. To determine if any pixel is green or not green, a simple analog comparison can be made between instantaneous R–Y and G–Y signals and also the instantaneous B–Y and G–Y signals. If in both cases the G–Y signal is greater, the pixel can be considered to be green.

In a preferred form of the invention the green state of a pixel is determined by operation of an algorithm wherein a pixel is deemed to be green when both of G>R and B<a set threshold for the blue component applies. This algorithm is preferable to the G>R and G>B algorithm above when the electronics to implement it is likely to be noisy and false green decisions are being returned. This is useful in low light conditions when present commercially available CCD cameras are in use. In this situation there is a component of noise present on the camera output signals. It has been found better to compare the B-Y signal to fixed reference voltage slightly offset from the signal level for black. This yields much better noise immunity while still providing a valid implementation of the above algorithm, since for a 'green' pixel the R-Y and B-Y signals are generally below the black signal level.

The detector can be any camera generating a colour output and typically it can be based on use of solid state devices such as charge coupled devices (CCD). The intensity of light which the device is to work with can vary considerably in open conditions and performance is enhanced by use of a hood whose function is to smooth out any marked light variation.

The detector and control circuitry which is used in the present invention is ideally able to locate weeds against a

3

variety of backgrounds such as black basalt soils, red soils, bare ground, stubble covered ground, rough rocky ground, changing light conditions, etc. It is found that a solid state detector such as a CCD based detector is best operated slightly out of focus so as to avoid false triggers which may otherwise arise when traversing ground having varying characteristics.

The circuitry which operates on the detector's signal is preferably able to perform its analysis in a short time so as to better typical efficient travel times of an agricultural spray. This is more readily enabled at lower costs by means of analogue circuits for processing the detector output.

The detector of the invention is used to convert an image of an area which is covered by the spray to a signal stream containing data which is equivalent to a picture frame which, when a solid state device is used typically comprises an array of pixels. The Red (R), Green (G) and Blue (B) components (RGB) of each of the pixels can be operated on to establish the green state of each pixel. A decision to spray might be based on the green state of a set of particular adjoining pixels or alternately the total or summed green component of a set length of a number of successive scan lines can be determined as the basis of the decision. These operations can be performed using either of digital or analogue techniques, or a combination thereof. The final green state which is calculated, is to determine a result being either a spray on, or a spray off decision.

The implementation of the above might be by way of circuitry providing a largely hardware approach to the problem of when to activate a spray or it might involve operations performed largely within a processor which is programmed to perform the desired functions.

BRIEF DESCRIPTION OF THE DRAWINGS

To enable the invention to be more fully understood, various preferred embodiments of the invention will now be described with reference to the accompanying drawings, in which:

FIG. 1 is a schematic plan view of an agricultural sprayer fitted in accordance with the present invention;

FIGS. 2A, 2B and 2C are diagrammatic views of how the field of view of a sensor unit may be utilised to advantage in the invention;

FIG. 3 is a side view showing a spray nozzle spraying a weed detected by the sensor unit;

FIG. 4 is the diagram of a circuit which may be used in a controller in accordance with the present invention; and

FIG. 5 is a circuit diagram showing another form for the circuitry for a controller in accordance with the invention;

FIGS. 6 to 9 illustrate a decision making process as might be implemented to determine if a detector output contains a plant to be sprayed.

DETAILED DESCRIPTION

The agricultural sprayer 10 is typically comprised of an extended boom, or booms supporting a linear array, or arrays of spray heads therealong, which boom, or booms, is or are trailed by, or mounted on a tractor 11 or other like type prime mover. Boom 12 can be fitted with a plurality of spaced apart, individually operable, spray heads comprising spray nozzles 13, arrayed therealong and ideally at regularly spaced intervals. The spray nozzles 13 can be connected to one or more spray tanks such as spray tank 14 by suitable pipes, lines or conduits 15, either individually or off a manifold. The spray heads may be any of those known in the

4

art. A standard valve, as utilised in the agricultural spray field can provide the means whereby a single spray head is able to be selectively operated. Valve 16 selectively allows the flow of spray chemicals from piping 15 to the nozzles 13, each nozzle 13 being selectively operable by selective activation of its respective valve under control of a controller connected thereto typically via a selectively operable activator. This is ideally achieved by electrical means with the controller switching sprays on via use of solenoids which open selected valves in the supply line, or lines to activate their respective spray heads. All of these elements can be chosen from amongst a range of readily available, off the shelf lines which will be selected according to standard criteria known to those in the art.

A plurality of the detectors can be provided on the boom 12 of FIG. 1. They can be arrayed therealong so as to cover the width of ground spanned by the boom. The field of view of a single one of the detectors may be such as to cover the ground beneath a number of adjacent sprays so that a detector is not required for each spray head. As seen in FIG. 3 a detector, typically a CCD based type detector 17 can be mounted in a housing, enclosure or hood 18 which is open at its bottom and which is arranged to be passed over the surface 19, on which there may be weeds to be sprayed, as the tractor draws the boom thereover. The surface being treated will typically be a field being prepared for a new crop, the field being either cleared of the last crop or having a stubble thereon. The housing 18 can be an opaque hood which is ideally arranged so as to stop all direct light falling on the target area and that way causing deep shadows therein. The hood 18 acts to diffuse light in the target area, the light being that which passes under the hood, into the field of view of the detector 17.

When a CCD type detector 17 passes over bare soil or stubble, the CCD therein converts the image below into an output comprising a string of pixels each characterised by respective RGB components. The controller can then determine the greenness of each pixel by manipulations of its components. The signal which is output by the detector 17 can be examined to determine if the weed covers an area of greater than a preset size. If the green signal exceeds a preset threshold limit at which the spray is to be activated, the valve 16 can be activated to switch flow to the appropriate spray nozzle 13 to spray the weed 24 (see FIG. 3). The circuitry interconnecting the detector 17 and the nozzles 13 can incorporate a time delay so that the spray nozzle operates for a preset time so that all of a target weed's area is sprayed as the boom moves over it.

One CCD detector can run a number of spray heads, depending on the width of its viewing area, and generally four is typical. The distance from the camera to the ground is the factor which determines this. For example, if it is desired to use one camera to run six spray nozzles then the camera may be set higher to cover a greater area at the ground (see the comparison shown between FIGS. 2A and 2B). Alternatively it is possible to use a wider angle lens (comparison shown between FIGS. 2A and 2C). In reference to FIGS. 2A, 2B and 2C, 20 is the camera head, 21 is the viewing angle.

The selection of height of the camera and the lens characteristics will ideally be decided depending on what in field conditions the machine incorporating the controller is working with in working with a wheat stubble, an acute angle lens mounted higher will allow it to look more effectively down into the stubble whereas in the normal bare fallow, a wider angle lens could be used to look out further. The screening effect of stubble is enhanced as the viewing

angle decreases and the vertical stalks more effectively hide a small or flat weed not raised to the same degree above ground level.

The light diffusing hood's dimensions are not at all critical. The dimensions will be varied to allow it to be fitted to different booms. The hood is constructed and mounted to keep direct light from the viewing area.

If external lighting is to be used to allow night time operation, an even white light mounted in the light diffusing hood could be used.

Referring now to FIG. 4, the output from the CCD 123 is fed through an RGB decoder 140 and respective Red, Green and Blue digitizers 141–143 and then to a frame store 144. In the frame store the RGB components of the output of the CCD 123 can be stored in digital form. The information in the frame store 144 can be passed via RGB processor 145 to a Green discriminator 146 which monitors the level of the Green component using an algorithm such as the one described below in greater detail requiring both of G>R and G>B to exist in a pixel before it is deemed to be green with some consideration of the number of green pixels in an area before the decision is made to call the area in the field of view green and a weed. Alternately the algorithm which is operated can be G>R and B<a set value its described elsewhere herein. The discriminator 146 can operate a solenoid driver 147 which is operably connected to a valve associated with spray nozzle to activate it and spray the detected weed.

A size selection section can be employed. This size selection section can be used to check the number of green pixels in an area of the target area and if their number is above a preset threshold, it can activate the solenoid to control the flow of chemicals to the spray nozzle. The threshold could be made adjustable so that it can be varied to allow an operator to select the size of the plant to be detected.

The horizontal field of view of a detector can be divided into a number of smaller regions to allow a single detector and processing section to control multiple valves and associated sprays which can be activated by solenoids under control of the controller.

The digital circuit of FIG. 4 has two areas which add considerably to the cost and complexity. The first is that having the digitizers at the output of the detector means that the amount of data to be stored in the frame store for a frame of video data is high (of the order of 1 Mbyte) . The second is that in order to have a reasonable range of colour levels to process, 6 or 8 bit digitizers are required, which for video applications are rare and expensive.

In the embodiment of FIG. 5, the front end processing can be performed using analog componentry. In this case, only a 1 bit digitizer is required since the result of the comparison is either "green" or "not green". It should be noted that by using this analog implementation, the memory requirements in the frame store are eliminated and no expensive digitizers are required. The digital processing requirements are substantially reduced and the whole system speeded up.

Where determining the number of adjacent pixels digitally can be complex and expensive. A simpler and cheaper method to operate is one which counts the total number of green pixels in the horizontal lines instead of the number of adjacent green pixels and count adjacent vertical lines. FIG. 5 is a schematic illustrating the components of a circuit which can be used in the controller wherein an "is it green" algorithm is implemented at the front end. The detector 25 outputs its usual RGB components on respective lines 26, 27

and 28 respectively, connected in pairs to comparators with pair 26 and 27 fed to comparator 29 and 27 and 28 fed to comparator 30 which each produce a logic "1" (high) when the green component of the detector output is higher. The respective comparisons are examined by the AND circuit 31 and if both the comparators are logic "1" (high) ie, G>R and G>B, then a green signal, logic "1" (high) is passed to the one digitizer 32. The level of Green over Red and Blue can be made adjustable in the comparator circuits 29 and 30 by either enhancing the G signal or retarding the Red and Blue signals, so as to allow adjustment to take account of weeds with different green characteristics. If the comparator which determines G>B is disconnected from the green component in the detector output and its comparison is with a set value then the circuit will work with the algorithm requiring both of G>R and B<the set value to apply.

From the 1 bit digitizer the circuit feeds counters which may be ideally set up in a microprocessor under software control to implement the further processing of the detector output. The one bit digitizer increments either counter 33 or 34 depending on which region is being analyzed, with a programmable threshold therein, and if the number of green pixels in the line of the region being looked at exceeds this threshold then that line is considered green by storing a logic "1" in memory. Once all the lines in the region are analyzed and results stored, then the number of green lines are counted and these also have to exceed a preset threshold (Number) if a spray signal is to be generated. By using this two count method the width and height of a weed is determined. This reduces the amount of memory required while still providing similar results, at faster speed and as before the threshold can still be varied to allow selection of the plant size to be detected. For example, if the horizontal field of view of the camera is divided into four regions, the counting of the "green" pixels can be performed before any data is placed into the memory resulting in only 4 bits of data for each horizontal scan by the camera instead of perhaps 640 bits of data (80 bytes). This represents a reduction in the amount of data to be processed of over 90%.

The signal generated by the detector typically includes components for the three colours,' RGB, with each component characterised by both of hue and luminance. In the above set out front end algorithm, the RGB components can be the detector's values minus a factor which can be the luminance (Y) of the camera signal so as to work with pure colour signals. Depending on which camera is chosen, its output may be signals which are the equivalent of colour minus intensity. In the working with the signals R–Y, G–Y and B–Y, the controller is working with the pure colour components. These signal levels are normalised so as to produce more significant ratios at the comparators 29 and 30.

There are circumstances when the G>R together with G>B principle will break down.

Extreme intensity variations can adversely affect performance by making a CCD device for example underexpose or saturate. However, intensity variations can be smoothed out by use of the above described light diffusing hood.

In another circumstance, a specific gold colour has green higher than red even though it is not greenish. This problem might be overcome by seeing how close to G and R signals are and how close the G and B is. This is because the gold colour has a close G and R and nearly no blue.

In yet another circumstance, the CCD camera views dead (golden coloured) grass and sees the dark area in between the dead leaves with a green hue. This causes false triggers.

7

As the size of the dark areas are generally small, size adjustments could be used to cut them out. However, size adjustment would limit the effectiveness of the size selectability by which a minimum size of weed to be treated is set. Also, the size of the dark area varies with changes in brightness during the day. One solution of this problem is to vary the focus of the camera slightly off normal. This smears out these particular dark areas to cut the number of false triggers and they can be all but eliminated. The affect of focus could be reproduced within the electronics but as this increases complexity, it is best to work within the camera's focus. Focus is an analogue solution to a problem which might be worked digitally but at added cost.

In FIG. **5**, the circuit can account for when a "green" plant is straddling the boundary between two regions in the camera's horizontal field of view. Since it is customary to set up the spraying equipment to have an overlap region between adjacent spray nozzles, it is logical that an overlap region should also exist between adjacent regions in the "green" detection system. This can be performed as seen in FIG. **5**, by utilising two independent counters **33** and **34** to count the number of "green" pixels, and control when they start and stop so as to provide an overlap in the counting regions. This is seen in FIG. **5** wherein separate green pixel counters **33** and **34** are switched by a counter controller **35** and their total is compared with a threshold set by variable threshold **36**. The counters are synchronised so that counter **33** counts pixels in segment 1 (eg, pixels 0 to 140). Counter **34** counts pixels in segment 2 (eg, pixels 120–240). This gives an overlap at pixels 120 to 140 when a weed is straddling this area. Counter **33** then counts segment 3 whilst counter **34** counts segment 4. This is repeated through the range of pixels returned by the camera. Control counter **35** counts the range and resets the "green" counters **33** and **34**.

As stated above the examination of the detector output to determine the existence therein of a weed can involve, use of a microprocessor which performs the algorithm and establishes the green state of an area. FIGS. **6** to **9** show in flow chart form the sequence of operations by which a spray activation signal might be generated. This is illustrated with reference to the G>R and G>B version and area calculation based on a scan line approach.

FIG. **6** shows the main process operating with four regions (associated each with one of four spray heads). On start up at **150** the scan line process **151** (described below in greater detail with reference to FIG. **7**) is implemented. If the first region of a scan line is deemed to be green and the previous scan line was green in this region, see **153**, then counter is incremented at **154** otherwise it is cleared at **157** and the second region is processed (**158**) in the same manner. If the scan line counter for region **1** is incremented at **154** then the count is compared at **155** with a threshold and if it exceeds it then a solenoid on flag is set at **156** otherwise processing passes to region two. The forgoing processing is pursued through the third (**159**) and fourth (**160**) regions till the full frame is determined to be completed at **161**. At this point turn on and turn off times are set for solenoids whose flags are set and processing passes to the solenoid control process at **163** (described below in greater detail with reference to FIG. **9**.

The scan line process at **151** of FIG. **6** is seen in greater detail in FIG. **7**. On starting the scan line process at **164** the region process (described below in greater detail with reference to FIG. **8**) is implemented. If the last region on a scan line is determined to be processed at **166** then the scan line process exits to the is it green decision process at **152** of FIG. **6** otherwise the scan line process loops. The region process

8

at **165** is seen in FIG. **8** wherein on its commencement at **168** the detector output is examined pixel by pixel. On receipt of a pixel at **169** the algorithm G>R and G>B is implemented at **170**. If both conditions apply then a green pixel counter is incremented at **171** otherwise and the end of region is tested at **172** with processing looped to continue if the end of region is not reached. When it is processes continues with the green pixel count compared to a threshold at **173**. If the threshold is exceeded then a green region flag is set at **174** and processing passes back to the scan line process.

The solenoid control process is seen in greater detail in FIG. **9**. When the turn on and turn off times have been set for solenoids whose flags are set (see FIG. **6**) the solenoid control process is run. If a solenoid on state is indicated at **181** the solenoid is energised at **182** and so on through the set with this program exited at **183** and processing returning to the main process. At some cycle through the solenoid process a solenoid off state will be reached to signal that it is time to de-energise for any solenoid which is currently on.

As hereinbefore described, the circuitry preferably incorporates a time delay so that the spray nozzle will operate for a preset time after it activated. A timer circuit might be associated with the solenoid, holding it on for a preset time so that the activation signal need only be a switch on pulse. Alternately the activation signal might be held on for the requisite time.

Various changes and modifications may be made to the embodiments described and illustrated without departing from the invention as hereinafter set forth in the claims.

Some of the features of the invention may be summarised as follows.

The invention contemplates a first system for determining whether a pixel is to be deemed green, i.e:

to use the three R, G, B, signals from the camera (which are three voltages, or, if the camera has a digital output, three digital signals) directly in the algorithm, whereby the pixel is deemed "green" if, for the pixel: G>R and G >B

In another algorithm, the pixel is deemed "green" if, for the pixel: G>R and B <a predetermined value.

The invention also contemplates an alternative system for determining whether a pixel is to be deemed green, i.e:

the R, G, B signals from the camera are not used directly in the algorithm, but rather the R, G, and B signals are aggregated to produce a value for the light intensity (luminance, Y) according to the conventional formula:

$$Y=0.30*R+0.59*G+0.11*B$$

Thus, in the alternative, the algorithm for determining whether the pixel is or is not green is: the pixel is deemed "green" if, for the pixel: G–Y>R–Y and G–Y>B–Y.

The invention also contemplates the inclusion of a means for alleviating the effects of overexposure and underexposure of the scanned area.

When the areas of extreme light are infrequent, one solution is to activate the spray solenoids in these areas by default. The added security of ensuring that no "green" areas are missed is paid for with a slight increase in chemical usage.

To be able to discern these extreme light levels a signal known as "Luminance" is developed from the Red, Green and Blue signals from the camera. The signal in given as

$$Luminance=(0.3*Red)+(0.59*Green)+(0.11*Blue)$$

Luminance basically represents the image without any colour information, ie: it is what is viewed on a black and

5,924,239

9

white television or on a colour television if the colour
control is turned to its minimum position.

Once the luminance signal has been developed, the signal
level can be monitored for the extremes of either underex-
posure (dark areas) or overexposure (saturated light areas).
A: These conditions can then be used to either force the
system to regard them as "Green" areas and hence use
the same control mechanisms as are already present in
the system, or
B: Preferably, brought into separate counter system which
allow independent control of these conditions. This
added control allows the operator to decide whether to
conserve chemicals, or to ensure that no "green" areas
are left unsprayed at the expense of slightly higher
chemical usage.

2. PHYSICAL MEANS

Extremes of both underexposure (dark areas) and over-
exposure (saturated light areas) can be reduced to eliminate
default spraying with a corresponding reduction in chemical
usage to be fixing a light diffusing hood above the target
areas and keeps the target area/signal within the dynamic
range of the CCD. The reduced levels of ambient light have
no adverse effect as the electronic exposure control com-
pensates to match the light.

Patent publication DE-4,132,637 might be considered
relevant to the invention, in that it shows a (non-agricultural)
weed spray controller, in which weeds are detected by means
of a video signal.

We claim:

1. Method for the spot-application of a spray to green
weeds or other green plants in an agricultural field, charac-
terized by the following procedural steps:
viewing an area of ground with a colour sensing means,
the colour sensing means being effective to scan the
area in pixelated fashion and to issue three signals, in
respect of each pixel in turn of the scanned area, the
three signals being dependent, respectively, upon the
amount of Red, Green and Blue light reaching the
colour sensing means at that pixel;
comparing the Green signal of a pixel with the Red signal
and the Blue signal of the pixel, according to a prede-
termined algorithm relating the said three signals in
respect of each pixel of the scanned area and deeming
the pixel to have a "green" status or a "not green" status
in accordance with the comparison;
assimilating the statuses of the pixels in a patch of the
pixels, the extent of the patch being defined in that the
pixels making up the patch are linked to the other pixels
in the patch in accordance with a predetermined degree
of spacial and temporal proximity to each other within
the scanned area;
comparing the aggregate of statuses of the pixels of the
patch with a predetermined value, and of deeming the
status of the patch to be "green" or "not green" in
accordance with the comparison; and
in respect of each of a plurality of spray heads operating
a spray head to produce a pulse of spray over a patch
in accordance with the patch having the status of
"green".

2. Method of claim 1, wherein:
the method includes the step of computing the luminance
Y of the pixel; and
the Red, Green, and Blue signals as used in the algorithm
are R–Y, G–Y, and B–Y.

3. Method of claim 1, wherein the algorithm is of the form
in which the status of the pixel is set to "green" if both (a)

10

the Green signal exceeds the Red signal, and (b) the Green
signal exceeds the Blue signal.

4. Method of claim 1, wherein the algorithm is of the form
in which the status of the pixel is set to "green" if both (a)
the Green signal exceeds the Red signal, and (b) the Blue
signal is less than a predetermined value.

5. An agricultural spray controller by which to control
agricultural spray apparatus, which controller detects green
plants on a surface being treated to enable the spot appli-
cation thereto of a spray, comprising:
a surface viewing means for generating a pixelated colour
video output including red (R), green (G), and blue (B)
(RGB) colour components representing its field of
view; a pixel receiving means for determining the red,
green, and blue colour components of the pixels in the
video output;
a green-pixel-determining means for determining whether
each pixel is to be deemed green, depending on a
relationship of the colour components of the pixel;
a green-area-determining means for counting whether the
number of deemed-green pixels in the video output
corresponding to an area of the field of view exceeds a
pre-determined number, the predetermined number is
based on the number of deemed-green pixels deemed
indicative of the presence of a green plant in the area;
an output means for delivering a spray activation signal
responsive to the green-area-determining means when
the number of deemed-green pixels in the area exceeds
the predetermined number;
a first comparator to determine if G is greater than R;
a second comparator to determine if B is less than a set
value; and
a processor to produce an activation signal by which to
activate a spray head when G exceeds R and G exceeds
B, when B is below the set value.

6. An agricultural spray controller as claimed in claim 5
wherein the first and second comparators output to a digi-
tizer through an AND circuit, the digitizer out-putting a
green or not green state pixel by pixel to the processor.

7. An agricultural spray controller by which to control
agricultural spray apparatus, which controller detects green
plants on a surface being treated to enable the spot appli-
cation thereto of a spray, comprising:
a surface viewing means for generating a pixelated colour
video output including red (R), green (G), and blue (B)
(RGB) colour components representing its field of
view;
a pixel receiving means for determining the red, green,
and blue colour components of the pixels in the video
output;
a green-pixel-determining means for determining whether
each pixel is to be deemed green, depending on a
relationship of the colour components of the pixel;
a green-area-determining means for counting whether the
number of deemed-green pixels in the video output
corresponding to an area of the field of view exceeds a
pre-determined number, the predetermined number is
based on the number of deemed-green pixels deemed
indicative of the presence of a green plant in the area;
and
an output means for delivering a spray activation signal
responsive to the green-area-determining means when
the number of deemed-green pixels in the area exceeds
the predetermined number, wherein pixels are counted
across a scan line, segment by segment over the area of

the field of view and if the number of pixels deemed to be green in a line in a segment exceeds a threshold then the segment line is deemed green, the number of scan lines in a segment deemed green are counted and if the line count exceeds a threshold then a plant is deemed to exist in that segment of the field of view.

8. An agricultural spray controller by which to control agricultural spray apparatus, which controller detects green plants on a surface being treated to enable the spot application thereto of a spray, comprising:

a surface viewing means for generating a pixelated colour video output including red (R), green (G). and blue (B) (RGB) colour components representing its field of view;

a pixel receiving means for determining the red, green, and blue colour components of the pixels in the video output;

a green-pixel-determining means for determining whether each pixel is to be deemed green, depending on a relationship of the colour components of the pixel;

a green-area-determining means for counting whether the number of deemed-green pixels in the video output corresponding to an area of the field of view exceeds a pre-determined number, the predetermined number is based on the number of deemed-green pixels deemed indicative of the presence of a green plant in the area; and

an output means for delivering a spray activation signal responsive to the green-area-determining means when the number of deemed-green pixels in the area exceeds the predetermined number, wherein the RGB colour

components of a pixel are examined and if G>R and G>B then the pixel is deemed to be green.

9. An agricultural spray controller by which to control agricultural spray apparatus, which controller detects green plants on a surface being treated to enable the spot application thereto of a spray, comprising:

a surface viewing means for generating a pixelated colour video output including red (R), green (G), and blue (B) (RGB) colour components representing its field of view;

a pixel receiving means for determining the red, green, and blue colour components of the pixels in the video output:

a green-pixel-determining means for determining whether each pixel is to be deemed green, depending on a relationship of the colour components of the pixel;

a green-area-determining means for counting whether the number of deemed-green pixels in the video output corresponding to an area of the field of view exceeds a pre-determined number, the predetermined number is based on the number of deemed-green pixels deemed indicative of the presence of a green plant in the area; and

an output means for delivering a spray activation signal responsive to the green-area-determining means when the number of deemed-green pixels in the area exceeds the predetermined number, wherein the RGB colour components of a pixel are examined and if G>R and B<a set threshold for blue then the pixel is deemed to be green.

* * * * *

Appendix F

Real-time LBP implementation

```c
#define IMAGE_WIDTH 640

#define IMAGE_HEIGHT 480

#define GUIDANCE_STACKSIZE 307200

#define COMPONENT_MAX_SIZE 100000

uchar retainedImage[GUIDANCE_STACKSIZE];

int guidanceComponentStack[COMPONENT_MAX_SIZE];

int componentStackPtr=0;

const int widthMinusOne=IMAGE_WIDTH-1 ;

int CCPos=0;

int guidanceStackPointer=0;

int guidanceStack[GUIDANCE_STACKSIZE];

int guidanceHighSize=0;

const int positionOne=-1*IMAGE_WIDTH;

const int positionTwo=-1*(IMAGE_WIDTH-1);

const int positionThree= 1;

const int positionFour=IMAGE_WIDTH+1;

const int positionFive=IMAGE_WIDTH;

const int positionSix=IMAGE_WIDTH-1 ;

const int positionSeven= -1;

const int positionEight=-1*(IMAGE_WIDTH+1) ;


void guidanceLBPFunction(int xstart,int xstop,int ystart,int ystop, IplImage*
grayScaleSourceImage, IplImage*returnedImage)
{
        //this is a 3x3  lbp

        uchar* gray=(uchar*)grayScaleSourceImage->imageData;;

        uchar* returnedImg=(uchar*)returnedImage->imageData;

        uchar* currentPos;

        int lbp_data[9];

        int loop=0;

        int centre_val=0;

        int pVal=0;
```

```cpp
const int noiseOffset=2;// this value is subtracted from the current positions
//value so that there needs to be a reasonable change to find edge
int currentPositionOffset=0;
for(int x=0;x<IMAGE_WIDTH;x){
        for(int y=0;y<IMAGE_HEIGHT;y){
                currentPositionOffset=y*IMAGE_WIDTH+x;
                if(x>=xstart&&x<xstop&&y>=ystart&&y<ystop){
                        currentPos=gray+currentPositionOffset;
                        centre_val=*(currentPos)-noiseOffset;
                        loop=0;
                        pVal=0;
                        if(centre_val>*(currentPos+positionOne))lbp_data[loop++]=0;
                        else{lbp_data[loop++]=1;}// 1

                        if(centre_val>*(currentPos+positionTwo))lbp_data[loop++]=0;
                        else {lbp_data[loop++]=1;}//2

                        if(centre_val>*(currentPos+positionThree))lbp_data[loop++]=0;
                        else lbp_data[loop++]=1;//3

                        if(centre_val>*(currentPos+positionFour))lbp_data[loop++]=0;
                        else {lbp_data[loop++]=1;}//4

                        if(centre_val>*(currentPos+positionFive))lbp_data[loop++]=0;
                        else {lbp_data[loop++]=1;}//5

                        if(centre_val>*(currentPos+positionSix))lbp_data[loop++]=0;
                        else {lbp_data[loop++]=1;}//6

                        if(centre_val>*(currentPos+positionSeven))lbp_data[loop++]=0;
                        else lbp_data[loop++]=1;//7
```

```
                if(centre_val>*(currentPos+positionEight))lbp_data[loop++]=0;

                else {lbp_data[loop++]=1;}//8


                //determin if uniform or non-uniform and what value if uniform

                int transitions_high=0;

                int num_high=0;

                int once_high=0;

                int loop_number=0;

                int contigousStart=0;

                for(int looping=0;looping<8;looping++){

                        if(looping==0&&lbp_data[looping]==1)contigousStart=1;

                        if(lbp_data[looping]==1&&!once_high){

                                transitions_high++;

                                num_high++;

                                once_high=1;

                        }

                        else if(lbp_data[looping]==1&&once_high){

                                num_high++;

                        }

                        else once_high=0;

                }

        //if 1 was high and 8 was high and only had one transition from low to high in the circle

                if(contigousStart==1&&once_high==1&&transitions_high==2)

                        pVal=num_high;

                else if(!transitions_high)pVal=0;//blank pixels- all 0 intensity
values

                else if(transitions_high==1){//all other uniform results

                        pVal=num_high;

                }

                else pVal=9;// non-uniform

                *(returnedImg+currentPositionOffset)=pVal;//*20;// the *20 is so
that the image is easier to view  by eye.- not needed for the processor
```

```
            }

            else *(returnedImg+currentPositionOffset)=0;// clear all values
outside of ROI

            y++;

        }

        x++;

    }

}
```

Appendix G

Side-shift hitch operation for row guidance

A side shift hitch is a tool for implement guidance. The side shift hitch takes control signals from sensor systems (vision system in the pyrethrum application) and corrects the implement position in the row, based upon the guidance algorithms in the sensor system.

The side shift hitch is sandwiched between the tractor and the implement. The hitch is in two sections. One section is mounted rigidly on the tractor and the second section is mounted on the front of the implement. The second section can slide left or right within the first section. The sliding is achieved by means of a hydraulic ram triggered from a row detection system. In this case the row detection system is a camera system (the same camera as the spot spray camera). Figure G.1 is an image of a side shift hitch highlighted in a red ellipse, in position between the tractor and the implement. Figure G.2 is a drawing representation of how a side shift hitch operates.

Figure G.1: Image of side shift hitch (in the red ellipse) in position between the tractor at the front and the implement at the rear.
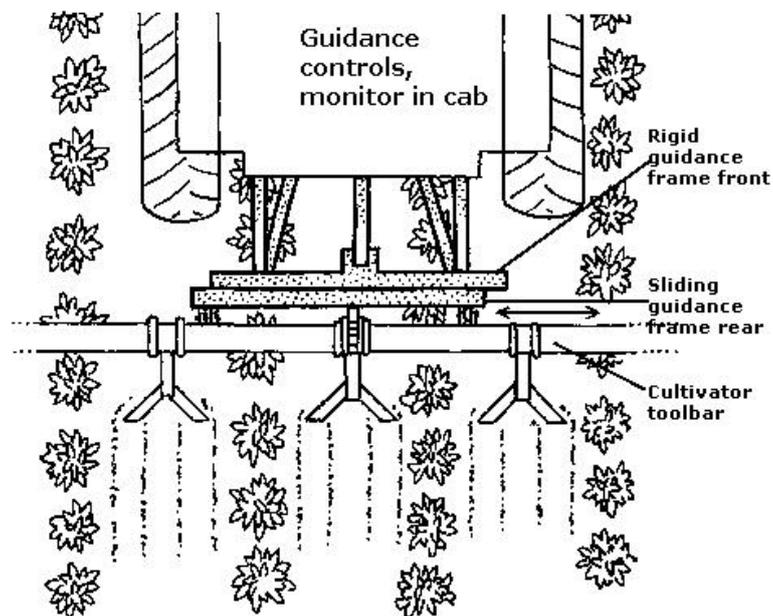

Figure G.2: Drawing of a side shift hitch components reproduced from http://www.sare.org.