Loughborough University
Institutional Repository

Plant classification combining colour and spectral cameras for weed control purposes

This item was submitted to Loughborough University's Institutional Repository by the/an author.

Additional Information:

- A doctoral thesis submitted in partial fulfilment of the requirements for the award of Doctor of Philosophy of Loughborough University.

Metadata Record: https://dspace.lboro.ac.uk/2134/4078

Publisher: © Pauli J. Komi

Please cite the published version.

PLANT CLASSIFICATION COMBINING COLOUR AND SPECTRAL CAMERAS FOR WEED CONTROL PURPOSES

by

Pauli J. Komi

A thesis submitted in partial fulfilment of the requirements for the award of

Doctor of Philosophy

Loughborough University

2008

CONTENTS

LIST OF FIGURES

PLANT CLASSIFICATION COMBINING COLOUR AND SPECTRAL CAMERAS FOR WEED CONTROL PURPOSES

ABSTRACT

This study was conducted to design and evaluate a novel dual camera sensor for use in an accurate single leaf level plant detection and classification system for weed control purposes. The system was to utilise and combine the benefits of colour and spectral imaging technologies together with novel data processing techniques. Such combination of colour and spectral imaging devices has not been previously used in precision agriculture.

Environmental consciousness and requirements for production volumes of organic produce are constantly increasing. Reductions or total elimination of chemical spraying is needed, and a technological solution of automating the weed control has been seen as one solution to solve the limitations in current crop production methods. Recent studies have shown automatic plant detection and classification to be the only economically viable solution for the problem of automatic weed control.

Previous detection systems have shown adequate capabilities to detect and classify weeds and crop plants with certain limitations. Depending on the system, these limitations have been in spatial accuracy, operation in certain lighting conditions or selection of plants to be classified. A flexible system capable of robust plant classification under any circumstances and plant combinations has not yet been realised. It would be desirable to introduce a system capable of detecting any plant species and plants separately, thus allowing targeted and optimal weed control methods for each plant species.

The proposed system addressed the problem of automatic plant detection and classification by providing sub-centimetre level information on plant part locations separately for each plant species. This information could then be directly used to guide mechanical weeding tools or precision sprayers. The detection system was based on a novel combination of a sub-millimetre level colour camera and an accurate hyperspectral line scanning camera (spectrometer) in the spectral range of $400 - 1000$ nm. The spatial accuracy of the spectrometer was approximately five times lower than that of the colour camera. The system operated under controlled lighting conditions.

The colour camera allowed precise segmentation of plant borders, while the spectral camera produced detailed reflectance information to discriminate between plant types. The

system was able to collect data for classification from an area on a plant of approximately 6.5 mm by 6.5 mm, although typically areas as small as 3.5 mm by 3.5 mm were detected. These were also the spatial resolution boundaries of the system with the used test settings.

The system was first designed and evaluated in laboratory conditions using controlled lighting and a selection of leaves from 6 plants. Data collection and analysis methods were designed for a scanning system with simultaneous image acquisition from both cameras. Shape, colour and spectral reflectance information were used to correctly classify these individual leaves with a probability of up to 98% using linear stepwise discriminant analysis.

A method of classifying separate leaves is not robust in a real field environment where plant leaves are often overlapping. This makes the use of shape calculations difficult. A novel method of extracting data from the small windows was proposed. Colour and spectral data within these windows was classified separately and the windows formed a grid like structure with approximately 3.5 mm spacing between them. This allowed spatial filtering of the classification data and noise reduction for the results by utilising information in a 3 by 3 window neighbourhood around each data window. During laboratory tests the windows for 6 plant leaves were correctly classified at 97.8% when the leaves were separated, and at 85.2% with overlapping leaves.

The system operation was also evaluated in real field conditions. Four crop plants and 16 weed plant types were imaged on a field over a period of 11 to 25 days after sowing. Total average classification performance in field conditions with linear discriminant analysis was up to 85.1%, while classification results investigated as a two-class case of crop vs. weed plants was up to 99.5% and 83.8%, respectively. The spatial filtering method was shown to improve results on average by 7.5%.

Plant reflectance measurements on different days allowed a novel analysis of short term temporal changes in the plant spectra due to growing conditions and growth stages. Analysis on short term spectral changes and their effects on classification accuracies have not been found in the literature. The temporal analysis showed that the average spectra of any plant type changes considerably over a period of just few days, and has a trend like behaviour when investigated at individual wavelengths. Classification models with training set data from previous days did not perform well. This indicates the need to have an up to date training set available at all times explaining the subtleties in local conditions.

The proposed detection and classification system with intelligent data processing methods has been shown to perform at a comparable level with previous systems. The novel system does not suffer from the typical limitations of previous systems, and is

flexible to be used with any plant types in their early growth stages. There is also potential to include plant height, shape or any other relevant feature to the classification for increased robustness.

The presented data processing method allows considerable processing and data reductions within the camera hardware. Only small fractions of the processed image data would need to be transferred via the camera interface. This would compensate for the increased data flow created by using two cameras. Therefore, the real-time implementation of the system is thought possible with the right hardware choices and optimised data processing algorithms.

ACKNOWLEDGEMENTS

Thank you, professors Mike Jackson and Rob Parkin for making my PhD project possible and for the support you have given throughout the project.

Thank you Dave Liddell for your help with technical problems and a special thanks to Terry Harrision for all the guidance and entertainment as long as a cup of tea was waiting. I would also like to thank the green of the golf courses and the snooker table to take my mind off the work when most needed.

Finally, I offer my greatest gratitude for my wife Erin for the continuous love and support during my project. I cannot express how valuable the hours you have spent thinking of my work on top of everything else have been. Also, I am grateful for the added motivation in the last stages of the writeup. I look forward to the next steps in our lives - together as a family.

1.0   INTRODUCTION

Modern agriculture is becoming more reliant on computer-based systems. Various technical advances have opened new possibilities to gather information and utilise it in agriculture as well as in other disciplines. Agriculture may not have traditionally been the first to implement the latest discoveries in technology, however, precision agriculture with localisation such as Global Positioning System (GPS) and other information technologies are becoming every day tools for farmers. Automated machines are starting to take over tedious tasks formerly performed only by humans.

Economic and ecological benefits are the driving forces to implement new methods into agriculture. Balancing efficient farming and preservation of nature has traditionally been difficult. Technologies for agricultural tasks help the farmers to log their actions, make better decisions and automate activities for faster and more economical execution, while saving energy and reducing the amount of chemicals used.

Agriculture has yet to see fully autonomous vehicles in wider commercial use, but time consuming and straining tasks have been the first to be allocated plenty of research resources. These tasks include cultivation, planting, weed control and harvesting. Environmentally conscious farming is rightfully in high demand and this is exactly where new technologies can offer solutions.

The current and ever-growing market for organic produce in the UK in 2004 was some £1.2 billion or 1.05% of the grocery market. The same figure in the whole world was £15.5 billion (SoilAssociation 2005). Organic food production can be promoted on the basis of several benefits; healthier food, improved farming environment and contribution to the rural economy. Organic farms have even out-performed the non-organic farms in mean sales values per hectare in one particular study (Lobley, Reed et al. 2005). New weed control methods and agricultural technologies are required to fulfill the demands of continuously increasing organic production volumes.

Precision agriculture is a term used for agricultural systems that exploit some of the new technological advancements in agricultural site-specific management. Applications of precision agriculture include variable-rate spraying, automated row following for harvesters, mapping of weed patches and crop yield management (Baerdemaeker, Munack et al. 2001). Minimising and making the use of chemicals as effective and non-disturbing as possible is the objective with many new technologies used in agriculture. Ultimately the goal could be to stop the use of chemicals completely.

According to Moran, Inoue et al. (1997) there are three basic types of information

required in precision agriculture. These are information on seasonally stable conditions (e.g. annual yield and soil based properties), information on seasonally variable conditions (e.g. weed infestations and weather) and information required to diagnose, find cause and develop strategies for crop management. An imaging system has potential to provide information for all three categories. The different variables that can be measured include water conditions in soil (Levitt, Simpson et al. 1990), crop yields (Leon, Shaw et al. 2003) and weed populations (Thorp and Tian 2004). This information can then be used to target soil and weeds locally, determine the additional nutrient needs and vary the seed rate depending on the soil conditions (Godwin and Miller 2003).

After an extensive literature survey, an agricultural task of autonomous weed control has been divided into three main problems. They are localisation and navigation of the machine or implement, detection and classification of plants and weed control methods, as shown in Figure 1. The figure also shows some examples of sub-tasks under the main three problems. The sub-tasks with the bold outline are the ones that this research concentrates on. Although the problem of building a working autonomous weed controlling device is not limited to the three problems presented, one must still solve at least one sub-task from each of the three problems for success.
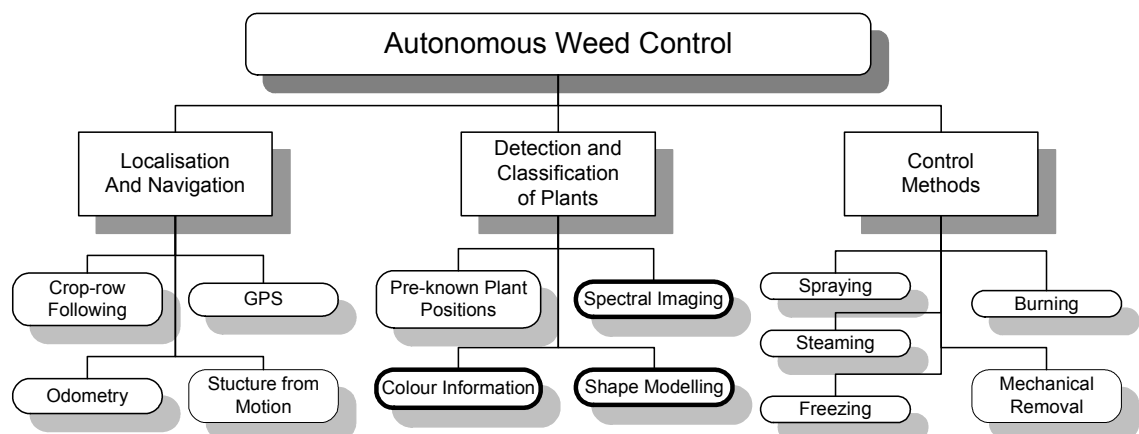
Figure 1 Examples of tasks involved in autonomous weed control

There are two different approaches to automatic weed control; mapping (off-line) and real-time (on-line) control. In the mapping approach the weed locations (or plant locations) are first mapped and then the spatial distribution information is used with an appropriate control method. This approach is illustrated in Figure 2 (a). Mapping of weeds is done prior to actual control actions and weed maps are used later on to apply appropriate control methods to specific locations. GPS position measurements can be used in mapping applications.

The real-time approach illustrated in Figure 2 (b) is more demanding in terms of computer processing power. Control methods need to be applied almost instantly after detection and there is limited time for measurement analysis. Operating speeds of many current systems are still below the economical threshold level due to the strict time limits in analysis. The most obvious benefits of the real-time approach over the mapping approach are in the potential to save time and energy as both detection and control are done simultaneously and the field is required to be covered only once. Control methods in both approaches can vary from uniform herbicide spraying to precise mechanical removal of weeds.
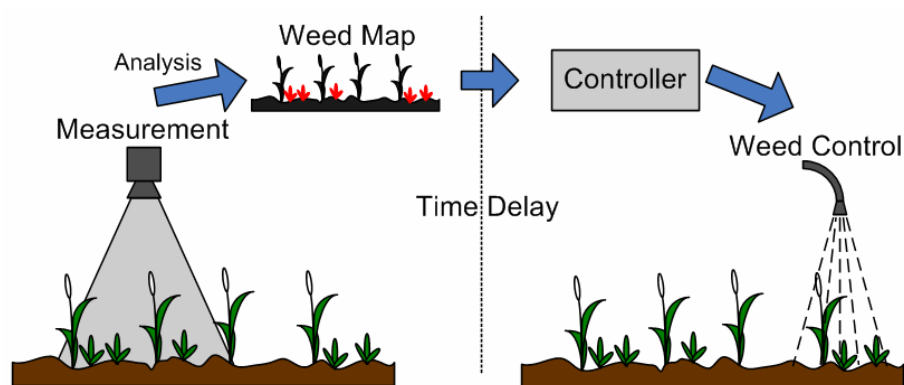
(a) Mapping approach in weed control
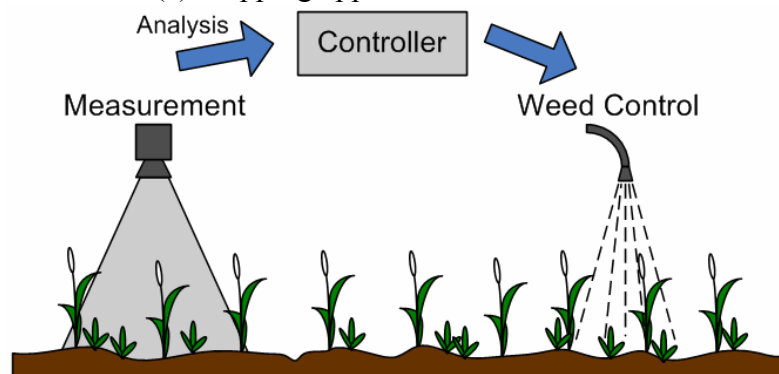
(b) Real-time approach in weed control. Reproduced from (Vrindts 2000)

Figure 2 Mapping and Real-time approaches in weed control

Weed detection methods almost always involve some kind of an optical system. Monochrome, colour and spectral cameras have been used to gather information for plant classification. Research projects until now have used a variety of spatial and spectral resolutions utilising only one imaging technology at a time. The common camera sensor technologies of Charged Coupled Device (CCD) and Complementary Metal–Oxide–Semiconductor (CMOS) are in use in almost all fields of imaging research. This also applies to agriculture, although applications may have their unique way of using a particular sensor. Passive and non-intrusive nature of imaging makes it suitable for numerous tasks.

Imaging is especially suitable for detecting weeds, and numerous methods are being used in image processing to analyse colour, shape and spectral information.

As seen in Figure 1, an automatic weed control system cannot operate without a solution to all three problems (navigation, weed detection and control). Mobile robots have great potential to offer a platform for such solutions. There is a vision that autonomous mobile robots will gradually take over tasks typically performed by humans and it has been suggested that mobile robots could be an economical solution for the particular application of autonomous weed control (Pedersen, Fountas et al. 2006).

There are numerous obstacles in creating autonomous robotic systems for agriculture. Harsh, ever changing outdoor weather conditions and uneven surfaces are just some examples of problems such systems will face.

Despite these obstacles, automation in agriculture is inevitable and desirable. Sensible automation can create immediate ecological and economical improvements in agriculture and crossing the conventional boundaries of scientific research areas is needed to achieve these improvements. The effects of successful research on agricultural automation have local and global effects, such as reduced chemical levels in food produce and increases in production volumes without sacrificing environmental values. These effects are well worth pursuing.

This project addresses the problem of automatic plant species detection and classification by presenting a novel dual camera system. The individual benefits of colour and spectral cameras were combined to allow sub-centimetre classification of small areas on plants. The proposed system was first designed for and tested in laboratory conditions by classifying individual plant leaves. High spatial resolution of the system allowed a unique method of processing the spectral and colour reflectance data in a grid formation and filtering the classification results without losing accuracy. The system operation in field conditions with several crop and weed plants was also tested successfully. During the fields tests the temporal changes in the reflectance properties were investigated. An analysis of local short term changes in plant reflectance properties and its effects to classification had not been conducted before.

In this thesis, Chapter 2 gives an introduction and an extensive literature review on the agricultural problem of weed control. The chapter addresses the weed control methods and the current research directions for automating weed detection. Emphasis is given to technologies and methods involved in optical recognition and classification. Examples of some relevant applications outside the area of agricultural engineering are also provided at

the end of the chapter.

Chapter 3 explains the proposed system and the adopted data processing methods. Image data collection, image processing and feature extraction using the dual camera system are presented. Two data processing methods are described; a full leaf method of extracting reflectance and shape data from a full leaf and the novel method of extracting classification features from small windows on leaves with the spatial post classification filtering of these windows for low noise results.

Chapter 4 shows the selection process of the system imaging and lighting hardware. The test rig design for laboratory measurements is also illustrated in this chapter.

Chapter 5 presents the leaf classification tests in laboratory conditions. Calibration of the test equipment is explained, the two data processing methods presented in Chapter 3 are applied to the measurements and the results shown and analysed.

Chapter 6 explains the final tests in real field conditions. The data collection and the modified processing methods are explained. The results are then presented with an extensive analysis, and comparisons of some classification methods are given.

Chapter 7 presents the conclusions of the project, and explains additional thoughts on the improvements for the presented detection and classification system together with further research opportunities and requirements for making the system work in real-time under minimal user input.

2.0 WEED CONTROL METHODS AND AUTOMATIC WEED DETECTION RESEARCH

A survey by the Food and Agriculture Organization (FAO) of the United Nations in the 1970s showed that more than one third of the potential world food harvest was destroyed by pests such as insects, diseases and weed. More recent studies have shown that the percentage is actually increasing. (Zimdahl 1993)

A weed plant can be described as a plant that is unwanted at a specific location at a given time. Farmers have fought against the weed populations for as long as land has been used for food production. In conventional agriculture this weed control contributes a considerable amount to the overall cost of the produce.

This chapter outlines the weed control methods and gives special attention in describing the current research in automating the weed detection and control. Automatic weed detection is the only viable solution for efficient reductions or exclusion of chemicals in crop production.

2.1 WEED CONTROL METHODS

Even a well maintained field may have around 20,000 weed seeds/m^2 in ploughing depth. Depending on the conditions the seeds may be able to germinate even 5-10 years later. (Agronet 2005) These figures alone can give an idea of the magnitude of the weed control problem.

Unfortunately, a complete solution to the weed control problems does not exist. Remarks on weed control methods made by Lötjönen, Jalli et al. (2002) in their literature review on perennial weeds in spring cereal production can be generalized for any crop production and weed control method. They note that timing of the control methods and knowledge of the life cycle of weeds is particularly important for successful results. For example mechanical weeding or application of herbicides at the wrong time may have no effect at all. In the worst case, the timing or wrong methods will do further damage to the produce. Then again, well timed mechanical, chemical or other control methods have great effects. It has been shown that good weed control within the first four to six weeks after sowing are critical to avoid yield reductions from weeds (Loux, Stachler et al. 2007). Figure 3 shows an example how this relates to the growing stages of wheat. The figure also shows the usual timing of chemical control where most post-emergence herbicides are recommended to be used and are at their most effective.

Approximate Wheat growth stages

Days	6-7	11-13		16-18	21-23	25-27		28-31	38	45	49
Weeks	1		2		3			4	5	6	7
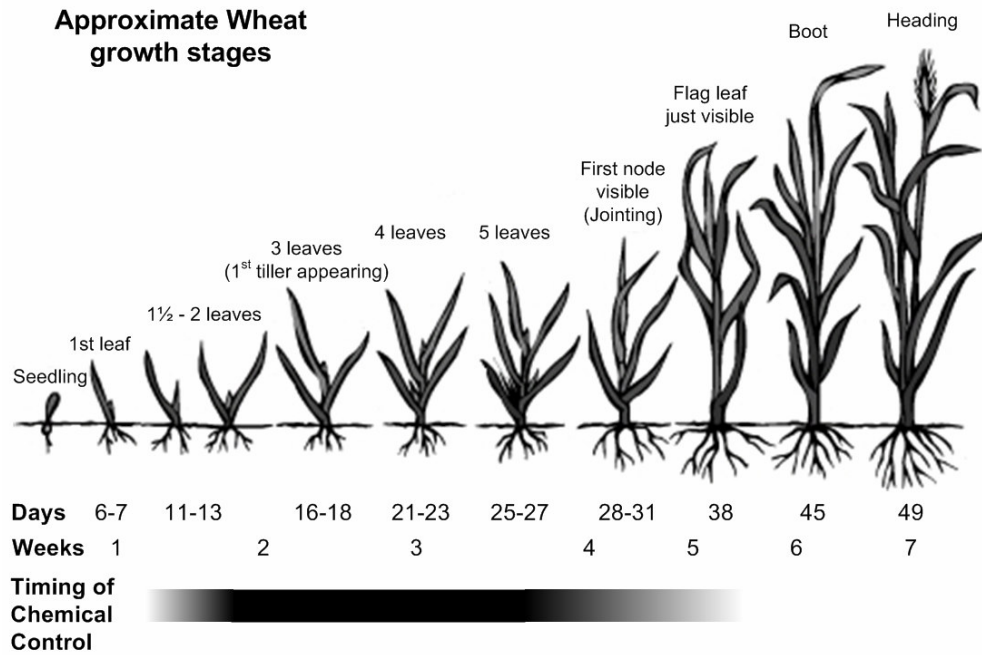
Figure 3 Approximate wheat growth stages and optimal timing of chemical control. Reproduced from Gregoire, Endres et al.(1997)

Three weeding zones have been identified in typical row sown crops. They are between crop rows (inter-row/between-row), along the crop rows between plants (intra-row/within-row) and close-to-crop (Blackmore 2004). A boundary of a few centimetres around the crop plant is typically considered to belong to the close-to-crop area. There are plenty of relatively simple solutions for between-row weeding in large fields. Currently the most common way is to pull a mechanical device in the back of a tractor that has gaps for crop rows. The results are then proportional to the driving accuracy of the operator. Automatic guidance of these between-row weeders can be done with good results as explained later in this section.

Between-row weeding is a significantly more complicated task. Any mechanical methods must avoid damaging the crop and crop detection methods need to be used. This is where most research on automating the weeding process is at the moment. Most current mechanical methods applied close to the crop will damage the crop and result in yield losses. None of the current weed control methods, apart from chemical and biological treatment, can deal with the weeds in the close-to-crop region.

Recently farmers and scientists have realized that attempts in removing all the weeds and exposing the bare land results in important disadvantages such as erosion, water run-off and nitrogen losses. Up to now, farmers have been most concerned on economical issues with chemical applications and have created resource allocation threshold levels for their use. The concerns on economical threshold could well be surpassed by ecological

concerns in the future (Hurle 1997). The right combination of economical and ecological thinking in weed control would be beneficial for long term economic sustainability without hurting the shorter term profit.

The best results in weed control are often achieved with the right combination of methods for the particular purpose. The following section will give an overview of common ways to control weed populations.

2.1.1 MECHANICAL WEED CONTROL

Weed treatment manually and with devices such as the hoe is considered mechanical weed control. Mechanical methods still dominate the weed control for many crops. Pulling weeds by hands, the most labour intensive method, works well for annual weeds, but perennial weeds typically reproduce and separate from roots before hand pulling is done. Hand hoeing is a more efficient and powerful way for weed control but it is time consuming (Zimdahl 1993).

Tilling is the common mechanical weed control method. During tilling, weeds are mostly buried by moving soil (Duval 1997). Other effects are separation of shoots from roots, stimulating the germination of seeds to be controlled by another tillage and exhaustion of carbohydrate reserves (Zimdahl 1993). Weed control is not the only reason for tillage since at the same time the compacted soil is broken up, aerated and cultivated.

Tilling has some obvious trade-offs. The main disadvantage of tillage is that although it breaks up the top layer, it generally compacts the soil and leaves it more vulnerable to erosion. It is also expensive to operate heavy machinery for long periods of time. It has been shown that reduced tillage can even have positive effects on weed control and yields in certain crops due to reduced disturbance to the crop (Blaise 2006).

The most common mechanical weed control devices are the rotary hoe and finger weeder illustrated in Figure 4. These machines can be operated at velocities between 6-20 km/h depending on the crop growth stage and the machine used.

Figure 4 (a) Rotary hoe (Singh 2006) and (b) finger weeder (Martens and Martens 2005) are the most common mechanical weeding devices

Most crops are grown in rows and uniform tillage is not recommended after about 30 days after sowing (DAS). However, between-row cultivation can still be done. The weeder needs to have gaps to match the crop rows and the operator needs to follow the rows precisely. Skilled operators can till within three centimetres of the row centres, but keeping high attention levels constant is a strenuous task. Automatic systems to follow crop rows have been created and are explained in §2.4.2.

2.1.2 NON-MECHANICAL WEED CONTROL

Non-mechanical weed control methods include thermal treatments, sound (Ultrahigh frequencies) and mulching (covering the soil) and exclude chemical treatments. Basically all the non-mechanical methods rely on thermal effects. These methods can be used in organic production instead of the chemical treatments. The thermal treatments can include burning and steaming of soil and plants. These methods are non-selective towards plant types and treat all in the same fashion. Intense heat stops the photosynthesis, cells can be dehydrated or plants are simply burnt away. (Zimdahl 1993)

Steaming can be costly and produce high the pollution levels. However, Danish researchers have developed an between-row band steamer that is commercially, economically and technically viable (Kristensen, Jørgensen et al. 2005; Sørensen and Jørgensen 2005). Their steamer treats all plants and seeds between the crop rows while keeping the temperatures near the crop rows low enough for the crop to survive. Compared to manual weeding, the cost reductions with band steaming were estimated at 72% and at 80% with robotic automatic steaming. Profitability of the robotic system is affected by potentially insufficient weed removal and high initial price of the equipment.

2.1.3 CHEMICAL WEED CONTROL

The spraying of herbicides is the most common and powerful weed control method at

the moment. Comparing a case without weed control to a case with chemical control, the total profit from yield including labour, machinery and herbicide costs can increase by about 70% when herbicides are used.

The economical benefits are clear, but some argue that the ecological effects overweigh the economical ones. All herbicides have some measurable toxicity levels to mammals and the life-cycle of most herbicides exceeds the plant life cycle resulting in herbicide levels carried forward from one crop season to the next (Zimdahl 1993).

Extensive use of chemicals in agricultural production has also led to problems with herbicide-resistant weed populations. For example, a study conducted in Western Australian fields showed that approximately 50% of the weeds were resistant to at least one type of the commonly used herbicides (Llewellyn and Powles 2001).

Many of the recent research projects in herbicide use have concentrated on reducing the amount of applied herbicides and targeting the spraying specifically to the weeds only. Compared to the conventional method of spraying everything (broadcast spraying or blanket spraying), cost and herbicide reductions of up to 80% are reported with systems detecting green vegetation (Antuniassi, Nery et al. 2003) or 11-90% with systems using pre made maps of weed coverage (Gerhards and Christensen 2003).

A four year study on weed populations in four different crops under precision spraying was conducted in Germany during 1996-2000 (Gerhards, Sökefeld et al. 2002). Maps of weed patches were created prior to spraying and DGPS (Differential Global Position System) signals were used to open spraying valves at locations where weeds had been located. Although the behaviour of the weed populations are not well known, the results clearly showed that great reductions can be made by precision spraying. A variable spraying system PatchSpray developed by Silsoe Research Institute in UK and commercialized by Micron uses premade weed maps to control the amount of chemicals being sprayed in a grid of approximately 2 x 2 m (Micron 2008).

Søgaard and Lund (2006) have researched the accuracy of spraying micro-doses of chemicals on only the plants. This technology could potentially further reduce the chemical amounts significantly as it was shown that spraying was possible to targets at sub-centimetre level.

In cases where weeds grow close to the crop, mechanical and non-mechanical control methods are often ineffective or more harmful than beneficial. Chemical or biological control methods may then be the only economically feasible ones.

2.1.4 BIOLOGICAL WEED CONTROL

Since weed populations resistant to herbicides have emerged, methods other than spraying are needed to fight these populations. In organic production, spraying is not an option. When mechanical methods are not suitable, biological ones may be considered. Biological weed control uses plants' natural enemies and plant pathogens (agent organisms) to target the specific plants. Organisms can be other plants, insects or fungi that typically target a narrow spectrum of weeds. They are best used in situations where typical chemical herbicides do not work well. Close to crop and chemically resistant weeds are examples of such cases. (Boyetchko 1997)

Biological control methods are still very much in the developing stages. The objectives are to maximize the impact on weeds and minimize the effect on crops, as with chemical treatment. Only a small number of bioherbicides have been registered so far (Hurle 1997).

Three main approaches for biological control have been identified: the inoculative or classical approach; the inundative or microbial herbicide approach; and the system management or augmentative approach. In the classical approach, exotic control organisms are introduced as many times as necessary. Microbial method is suitable for intensive (high volume) agriculture. The third approach, augmentative, aims at cautious manipulation of a weed-pathogen or weed-insect system (Müller-Schärer, Scheepens et al. 2000; Scheepens, Müller-Schärer et al. 2001).

Biological control methods do not typically produce immediate results like mechanical control. However, the results over time can be more advantageous. Biological weed control is not recommended to be used alone, but together with other methods it can prove to be very effective.

2.1.5 DISCUSSION

Certain weed control methods are clearly more suitable for the application of autonomous weed control than others. Chemical herbicide application is still the single most powerful control method, although many ethical and ecological arguments are against it. Precision spraying in patches or even on the single plant level is shown to reduce herbicide loads and the overall cost of spraying, and could be implemented in autonomous off-line or online applications.

Herbicides cannot be used in organic production and mechanical, non-mechanical or biological methods have to be used instead. In conventional crops grown in rows, there are no more problems in efficiently treating the weeds between the rows. For example, autonomous steaming and mechanical cultivation between rows has been shown to work

well in principle and practice.

The difficult locations for weeding are between the plants within crop rows. This is why an autonomous weeder could be equipped with a precise mechanical tool to disturb individual weed plants or with a precision sprayer. Even a combination of the two might prove effective.

In row sown crops the weed control methods in within-row areas could be applied with precision and some threshold for weed growth in between-row areas could be allowed for crops sown with large row spacing. Leaving some between-row weeds untouched may help reduce erosion and increase biodiversity. Finding the correct threshold such that crop plants would not have too much competition for nutrients etc. would be the key to this approach.

To apply the suggested control methods autonomously, a system capable of detecting specific weed locations has to be created. The next sections explain the requirements for such systems and present several methods in optical detection. The potential of these optical systems for the autonomous weeding task will be indicated.

2.2 AUTOMATIC PLANT DETECTION AND CLASSIFICATION

Automatic and efficient weed control in a real environment requires robust detection of crop and weed. There are several ways of gathering this information from manual human input to automatic recognition systems. For an economically viable system, automatic recognition and classification is the only solution.

The two automatic weed control approaches, as explained in Figure 2, are mapping and on-line methods. In mapping method the locations of weeds can be mapped for example with the help of a GPS signal and treated afterwards. In on-line method the detection and control happens seamlessly at the same time.

Weed detection and prevention can already begin before planting. Weed seeds could be recognized and removed, minimising the growth of unwanted plants (Granitto, Verdes et al. 2005). However, the remaining weed detection methods explained in this chapter are meant solely for the post-emergence stage.

Research on automatic detection systems started in the 1990's and was influenced by the need to create systems that would allow significant reductions in herbicide volumes. More or less all of the methods researched involved some form of digital optical system and analysis of the acquired data. Applying the knowledge of the detected weed locations has been extensively researched for use with precision chemical sprayers and mechanical methods.

Using optical imaging equipment in weed detection takes the research into the field of pattern or image recognition and classification. Four separate problems or steps can be isolated in this context. These problems are imaging, segmentation, feature extraction and classification. Each of the problems can and have been studied extensively and they are all needed for successful classification results. The solutions to these individual problems are generally application specific. This means that for example a classification system built for face detection does not work for plant classification and vice versa. Some important issues and previous research results regarding the specific application of plant detection and classification are listed in this chapter.

2.2.1 IMAGING

The first problem or set of problems to be solved for an automatic weed detection system is imaging. Aspects such as imaging platform setup, spatial resolution, spectral resolution and lighting conditions need to all be considered. The final design choice is highly dependent on the application and system requirements. The sub problem set for imaging is shown in Figure 5. The text in bold indicates the choices made in this project.
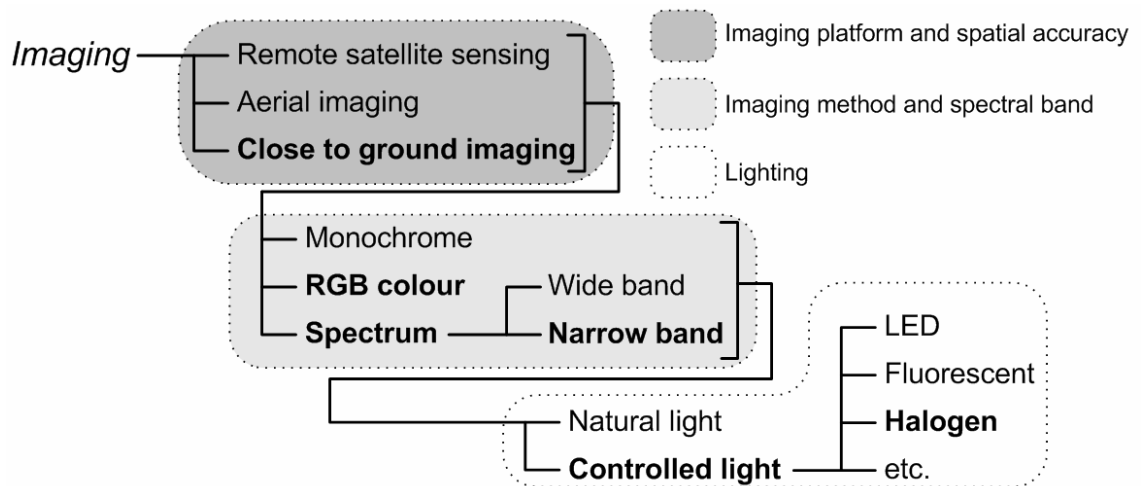
Figure 5 Issues involved with the imaging step of plant detection and classification

The first subset of imaging problems is related to the physical imaging platform and equipment. Remote sensing, aerial imaging and close to ground imaging are listed in this subset. All these imaging methods have been used in agricultural imaging tasks providing information at different spatial resolutions. This resolution is mainly dependent on the distance to the objects being imaged. The second subset of problems in imaging is a choice of spectral bands. This is closely related to the imaging platform and equipment selection.

Remote sensing is a term used for observations made at a considerable distance from

the target. Typically the term is used with satellite imaging, but aerial imaging can also be considered as a part of remote sensing. Information from satellites is commonly used for crop yield, weed population and insect infestation estimation at relatively low spatial resolution.

Nearly all satellite imaging applications include spectral imaging. The history of multispectral satellite sensing started with the launch of Landsat-1 in 1972. Until 1978 it gathered information from the earth under wavelength ranges from 500 to 1260 nm and from 500 to 750 nm with spatial accuracies of 75 m and 40 m, respectively. In one example the gathered multispectral data was used to identify 25 separate crops in an area of 458,000 acres in 45 hours. This and many more measurements have proven that remote sensing is a useful tool in agriculture (NASA 2005).

The spatial resolution of the early remote sensing satellites only allowed average plant canopy characteristics to be measured. Much higher spatial resolutions are needed for single plant level or even more accurate measurements. The latest satellites are capable of producing multispectral images with resolutions of up to just few meters (Moran, Inoue et al. 1997), while the price of remote sensing equipment has reduced dramatically. Smaller and more economical components for spectral imaging have become available mainly due to advances in digital imaging equipment (Gilchrist and Hyvärinen 2006).

Remote mapping of weed populations with aerial or land driven vehicles has been also shown to work well. The results are good enough for reducing the amount of herbicides used compared to broadcast spraying. Problems most often arise from mixed measurements where a gathered spectral data point can include information on different plants, soil residues or other objects. This is due to low spatial resolution of the data (Thorp and Tian 2004). An example on how high resolution colour images together with hyperspectral satellite remote sensing data can be utilised is given in §2.4.3.

Spatial resolution of an imaging system can be increased by increasing the imaging sensor resolution and the quality of the optics. Reducing the distance between the imaging device and the object is another obvious way of improving the spatial accuracy. To obtain single plant recognition levels the spatial resolution of a detection system needs to be at a sub-centimetre level.

With only a few exceptions, the remote sensing or aerial imaging devices image the ground from directly above (parallel to the ground). This is due to practical reasons in image registration and errors created in rectification of images taken at an angle. However, this arrangement may not always provide the best results when imaging close to the

ground. Selection of imaging angle can be important for certain applications. The difference the direction of the camera optical axis can make is illustrated in Figure 6 where the first image (a) is taken parallel to the ground and (b) at an angle. In (a) the position of the leaves on the ground can be easily detected, but single shoots or the crop rows of grassleaved plants are hard to detect. In (b) the crop rows can easily be detected while the exact location of the leaves on the ground is harder to measure.

(a)   (b)

Figure 6 Illustration of problems with plant detection. (a) Image taken parallel to the ground and (b) with a tilted camera showing the row structure (Andersen 2002).

Outdoor field conditions, where most weed detection and control takes place, further complicate the placement and use of the sensors. Dirt, changes in weather conditions and mechanical stress due to the vibration of ground vehicles introduce additional problems apart from complicated detection algorithms and technologies.

In plant detection the typical imaging devices are the RGB colour sensor and the spectral sensor. A monochrome detector simply does not provide enough information for plant discrimination and classification. Theoretically, the RGB sensor can be considered as a wide band spectral sensor with three bands, although usually only devices with several bands are considered spectral cameras. Spectral data adds a new dimension to the conventional evaluation of spatial brightness data from monochrome cameras by analysing the light intensity as a function of the wavelength. The terms multispectral and hyperspectral are indications of the spectral resolution of the system and are loosely used for systems with tens or hundreds of channels, respectively. The spectral resolution does not necessarily correlate with the spatial resolution.

Conventional CMOS and CCD imaging sensors are sensitive to the light in a spectral band of approximately 400 – 1000 nm. This band starts from the visible spectrum and ends at the near infra-red (NIR) range in the electromagnetic spectrum shown in Figure 7.
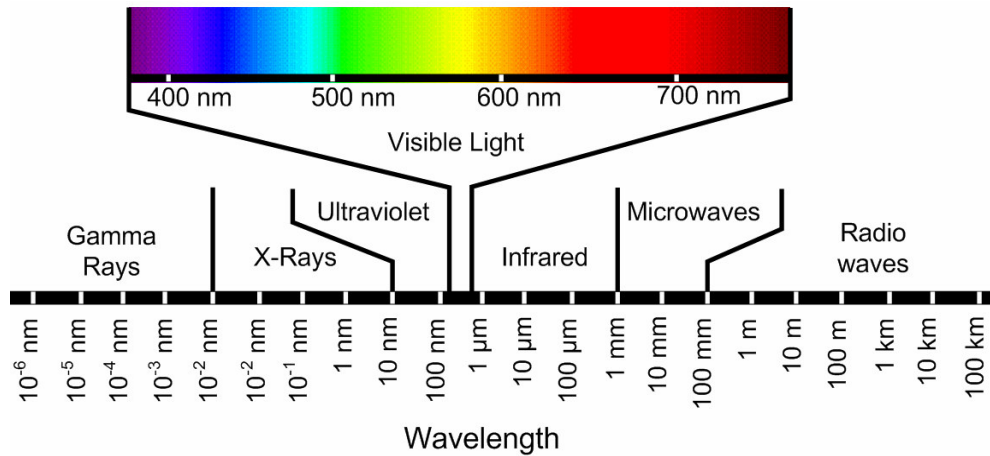
Figure 7 The electromagnetic spectrum

A typical spectral reflectance curve for a green leaf is shown in Figure 8. Perhaps the most important aspect of this reflectance curve is labelled in the figure as the 'red edge'. The term refers to an area in the spectrum where there is a rapid change in the reflectance of a living plant. This is a result of the absorption and reflectance properties of the chlorophyll in a leaf. Most of the light is absorbed below and reflected beyond 700 nm. The figure also indicates water absorption bands at higher wavelengths. These bands correspond to the natural frequencies of water molecules and the incoming radiation at these wavebands is absorbed in the watery leaf.
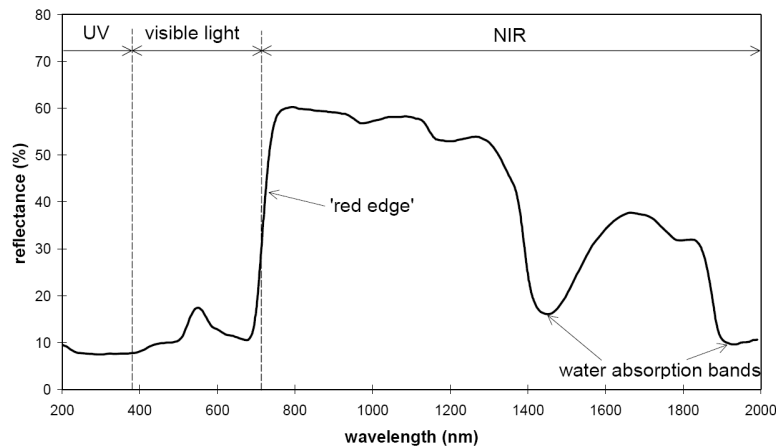
Figure 8 Typical spectral reflectance of a green leaf in ultra-violet (UV), visible (VIS) and near-infrared (NIR) wavebands (Vrindts 2000)

2.2.1.1 *Lighting Conditions in Imaging*

Most agricultural activities take place outdoors under uncontrolled lighting conditions. The sun may be the only source of light and the changes in the atmosphere, weather conditions and height of the sun affect the intensity and the angle of the light. Intensity variations are also caused by shadows. These changing lighting conditions can seriously affect the imaging parameters. One solution to minimise the changes in lighting conditions

is to block the sunlight and use artificial lighting i.e. to control the lighting. Controlled lighting eliminates the need to account for changing lighting conditions when image data is processed.

The effects of varying lighting conditions have been studied widely. One such study by Steward and Tian (1998) showed that simple soil and plant classification with RGB colour images was not greatly affected by changing lighting conditions. Another study by Steven (2004) concentrated on the effect of field of view in spectral measurements in varied lighting conditions. These results also showed that the soil and plant segmentation was not greatly affected by the lighting conditions and also that a mathematical relationship could be established between sunlit and shaded images.

Onyango and Marchant (2001) used a physics based model to adapt to changing lighting conditions with an RGB camera. A dichromatic reflection model and CIE daylight standard were used to investigate the performance of soil and plant segmentation under changing lighting conditions. Results were close to perfect but a limitation in their method is that the spectral irradiance in the light source needs to follow the CIE daylight standard.

A PhD thesis by Andersen (2002) concentrated on using computer vision in changing illumination conditions. His method of adaptation to the illumination changes was based on CIE daylight standard and dichromatic reflection models. The difference to the research by Onyango and Marchant (2001) is that Andersen investigated both an RGB camera and a three band camera with red, green and NIR channels. Andersen also compared results from controlled laboratory conditions to measurements from field conditions. He found that daylight could be modelled into two components; direct sun- and skylight. Using this model an adaptable daylight invariant soil and plant segmentation algorithm was developed.

Automatic gain control (AGC) is one way of trying to keep image properties somewhat constant in varying lighting conditions. This approach has been investigated with cameras with green, red and near-IR channels (Wright, Matthews et al. 2002; Xiang and Tian 2005). Especially the study by Xiang and Tian (2005) showed the potential of automating the gain control to create balanced images in varying lighting conditions. The approach was to use artificial neural networks and fuzzy logic to determine the desired gain values. AGC not only creates balanced images but also utilises the full dynamic range of the pixel depth in cameras.

Adaptation to changing lighting conditions has not been as widely studied with multispectral or hyperspectral systems. High resolution spectral data could potentially be

used directly to adapt to illumination changes. If a measurement of the current lighting conditions could be included in every measurement the adaptation should be trivial. In practice this would mean imaging a reference signal in every image frame or measuring the lighting conditions with an external device at all times.

Yao, Tian et al. (2002) have used an upwards pointing colorimeter in their outdoor field experiments to measure the illuminance and chromaticity data and used it as a reference signal to compensate for the illumination changes over time. In their experiment the main problem was the spectral mixing due to low spatial resolution spectral measurements. Therefore the results on adaptation to changing lighting conditions could not be isolated from the limitation due to spatial resolution.

Extra equipment, such as colorimeters or spectral measurement systems, adds to the critical total cost of the system. Therefore, a low-cost method could be considered and implemented. Placing a reflection reference plate permanently in the field of view of the spectral camera could be used to measure the current lighting conditions. Spectralon diffuse reflectance plate SRS-99-020 from LabSphere Inc., with a near uniform 90% reflectance over the spectral band of 400 – 1000 nm, is a possible reference plate option. The challenge of this method is the placement of such reference plate. Measuring the lighting conditions or a reference signal in one image location does not necessarily mean the lighting is uniform throughout the field of view. Changes may be introduced for example by shadows.

Lighting is an additional system and it is desirable to aim for agricultural systems with no controlled lighting. Even so, thinking of all the possible lighting conditions from complete darkness in the night time to bright sunshine during the day, controlled lighting is still the most desirable option in many cases. When planning to use controlled lighting it is important to choose the correct illumination conditions for the particular application. At best, suitable illumination can significantly reduce the required image processing. At worst, all of the desired characteristics cannot be detected due to the wrong wavelengths or direction of light.

Light source selection is especially important with spectral reflectance imaging. Different light sources radiate a different spectrum, and reflectance can only be measured at a particular waveband if some illumination on that band is available. Halogen lights provide a fairly smooth spectral curve for a broad wavelength range from visible to near infrared. Fluorescent lights have a few distinct spectral peaks in visible light and are generally suitable for RGB or monochrome imaging only. Generally speaking, bright

illumination enables shorter exposure times and the movement of the target is less likely to have an effect in the image.

2.2.1.2   Spectral Imaging Equipment

This section briefly describes the available spectral imaging technologies. The benefits and limitations of each technology are given.

Spectral imaging devices consist of two main components. First, a device is needed to separate the wavelength components in the spectrum; second, these components need to be sensed. The sensing is typically done using an imaging array. Wavelength separation is common with dispersive techniques using either a prism of diffraction grating. Common imaging arrays are CCD sensors for visible and near-infrared spectra (400 – 1000 nm) and inidium, gallium, arsenide (InGaAs) sensors for a range of 400 – 1700 nm (Kerekes 2006).

An imaging spectrometer gathers a three-dimensional data set, called a spectral cube. Two of the dimensions are spatial (x and y) and the third is the spectral wavelength. Various spectrograph technologies gather this spectral cube data in different modes. Figure 9 visualises how these modes are located in the spectral cube. A spectrometer with a grating and a linear one-dimensional array gathers data in a column (a), while a spectrally filtered two-dimensional array outputs a spatial plate with a thickness $\Delta\lambda$ specified by the filter properties (b). Finally, a spectrometer with a slit in the front and a grating and a two-dimensional imaging array gathers data about the plate (c).
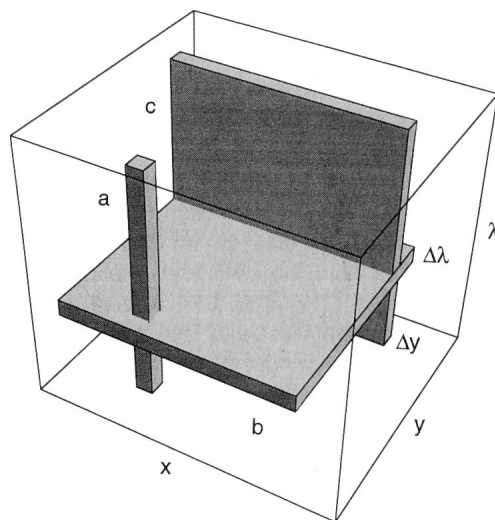
Figure 9 Spectrometer data gathering modes within a spectral data cube (Descour and Dereniak 1995)

When a device images only a portion of the spectral cube at once, a scan along the spectral or spatial axes is needed to gather the complete data in the spectral cube. This may require moving the spectrograph spatially, using a moving mirror, prism or grating setup or

changing spectral filter properties to scan all of the spectral components. Spectrographs with moving parts are not recommended to be used in harsh environments.

A spectrograph capable of imaging the whole spectral cube at once without moving parts has been presented by Descour and Dereniak (1995). This device uses three rotated gratings at 60° angles to create a set of dispersed images containing the spectral and spatial components. Estimation of the spectra requires heavy computations and in practice the system is limited in spectral and spatial accuracies such that its use for agricultural plant classification in large scale is impossible.

Another spectral imaging technique without moving parts uses a liquid crystal tunable filter (LCTF) to change the spectral properties of the system. The LCTF element can be placed in front of the objective lens or between the objective and the imaging sensor. The LCTF properties can be typically changed within $50 - 150$ ms and the possible spectral resolutions range from $0.25 - 20$ nm (CRI 2008). An image set taken with a set of filter properties constructs the complete data in the spectral cube. Examples of using an LCTF system in object recognition and classification are presented by Tominaga and Okajima (2000) and in wheat classification with images taken with a stationary camera by Alchanatis, Ridel et al. (2005). Systems with an LCTF can have excellent spatial and spectral resolutions, but switching time between the wavelength bands and the need to spatially register the different images limits their use in applications involving moving cameras or objects.

Spectrographs requiring spatial scanning can have high spectral and spatial resolutions. These systems are often used a in a pushbroom setup, where the equipment is scanning above the target, or with a mirror scanning the required field of view. Most aerial spectral imaging systems use the pushbroom setup. An example of this spectrograph type without moving parts is shown in Figure 10. In this figure, the spatial field of view is a narrow line specified by the objective lens and the entrance slit size. The spectral components within this line are dispersed to the imaging sensor through a prism-grating-prism system. The imaging sensor then creates an image with spatial and spectral axes. The presented Imspector spectrograph can be used to create a relatively low-cost and spectrally and spatially accurate imaging system capable of reliably imaging objects while moving at high speeds.
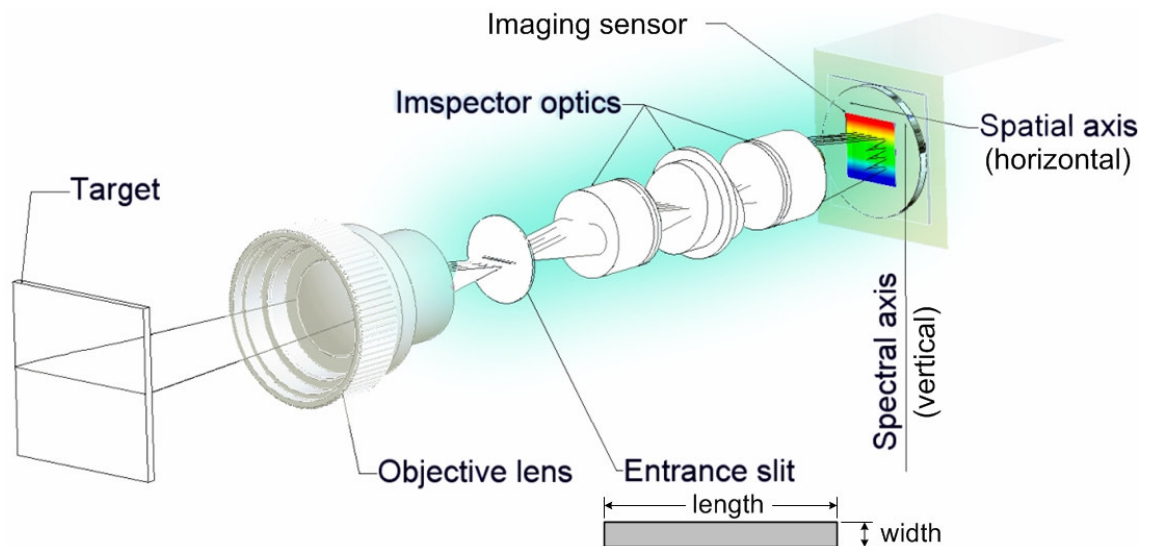
Figure 10 Imspector spectrograph functional diagram. Adapted from (Specim 2006)

2.2.2 SEGMENTATION

Segmentation process is the second step in the machine vision application of weed detection and classification. The purpose of this step is to segment or extract the plant and non-plant information from the images. Typically this means plant and soil segmentation. The most common methods of segmentation are listed in Figure 11.
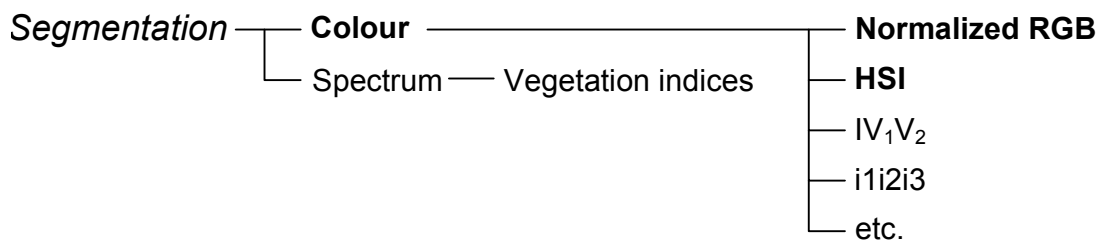
Figure 11 Methods of plant segmentation

Plant segmentation using RGB colour data is the most common method. The original RGB data is in three broad spectral bands of red, green and blue. Different RGB colour space transformations have been researched to find the one most suitable one for plant segmentation. A study by Phillipp and Rath (2002) compared discriminant analysis, canonical transformation, i1i2i3, Hue, Saturation and Intensity (HSI) and Lab conversions. A method based on discriminant analysis was found to perform best, but this method would require a lot of computing power and is not suitable for real-time processes. A slightly modified i1i2i3 colour space transformation i1i2i3$_{new}$ was suggested to be used instead.

Other segmentation methods include physics based dichromatic reflection models to segment RGB images to soil and vegetation (Onyango and Marchant 2001) and an IV_1V_2

transformation developed specifically for vegetation segmentation by Steward and Tian (1998). More traditional method of thresholding in HSI space was used by Hemming and Rath (2001). Their results showed that reasonably good segmentation was possible with this transformation.

Commercial products WeedSeeker and GreenSeeker from a Californian company (NTect Industries 2003) rely solely on detecting the infra-red reflectance of a plant canopy. These systems are not strictly speaking spectral measurement systems, although they work outside of the visible spectral range. The detection of plant material is based on relatively high near infrared reflectance of plants only. The separation to plants and non-plants is done by evaluating the reflectance value in infra-red. Once a set threshold in reflectance is reached within the field of view the image is considered to be of plant material.

When more than the three colour data channels are available the system can be considered spectral. Imaging systems with red, green and NIR channels and some similar combinations have also been developed to improve the segmentation as compared to the conventional RGB segmentation. The use of spectral data or the addition of a NIR band is based on using the physical reflectance properties of the chlorophyll content in the plants for segmentation. The most helpful property in the spectrum for segmentation purposes, the red edge, is illustrated in Figure 8.

A plethora of methods have been developed to utilize the red edge information in plant segmentation. The most common application is in satellite remote sensing and estimation of vegetation density. These methods are commonly called as vegetation indices. The vegetation indices basically estimate the amount of chlorophyll in the data. Some of these methods are listed in Huete, HuiQing et al. (1997), Thorp, Tian et al. (2002) and Thorp and Tian (2004). The most common vegetation index is the normalised difference vegetation index (NVDI), which is calculated as in the following equation, where NIR represents the near infrared and VIS the visible spectrum.

$$NVDI = \frac{\lambda_{NIR} - \lambda_{VIS}}{\lambda_{NIR} + \lambda_{VIS}} \qquad (2.1)$$

There is great potential for using spectral data in segmentation but often due to the lower spatial resolution it suffers from mixed signals. A mixed signal contains combined spectra from two or more objects when the field of view of a spectral data point covers an area with multiple reflectance properties. Colour data typically has a higher spatial resolution and does not suffer from similar mixing problems.

2.2.3 PLANT RECOGNITION

The next step after plant segmentation in the application of weed recognition and classification is plant recognition. During the recognition process, the segmented areas are separated into individual plants. These recognised areas are then used as inputs to the next step in the process, the feature extraction. The recognition step and some of the methods involved in it are listed in Figure 12. Sometimes the methods used in recognition cannot be separated from the ones used in feature extraction or even classification. One method may give a result to all these three steps at once.

Recognition — **Leaves / plant parts** — Wateshed
— **Whole plants** — Active shape models — Erosion / dilation
— **Whole blob detection** — Deformable templates
— etc. — Cutting and merging
— etc.
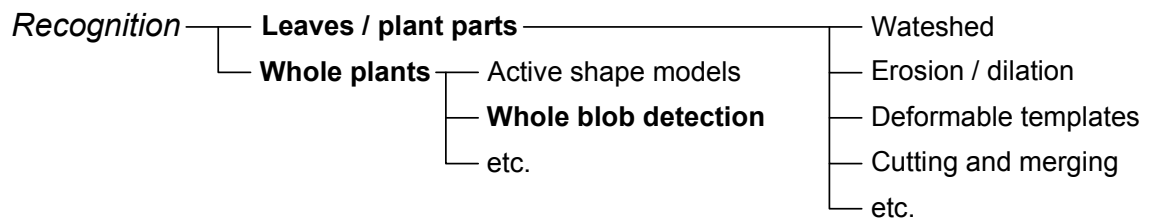
Figure 12 Recognition step in weed detection and classification

Recognition generally deals directly with segmented binary images. Individual regions in the segmented image, also known as blobs, may consist of overlapping plant parts and are rarely single and well defined leaves in the real field environment. The problem of recognising individual plants from the binary images can be easily understood by investigating a sample image of the cabbage plants and weeds shown in Figure 13. Many of the plant leaves between individual plants are connected to each other and would appear in the same blob after simple segmentation. One of the biggest problems in plant recognition in real environment is the unknown nature of shape and location of living plants. The overlapping plant parts are difficult for machine vision systems to deal with.

Figure 13 Example image of cabbage plants and various weeds
(Hemming and Rath 2001)

Recognition can divided into two main approaches; full plant/full blob approach and

an approach where an attempt is made to identify individual and potentially different plant parts. Recognising only full blobs is trivial and requires no additional intelligence but may not work well with real world field data that contains connected plant parts.

Active shape models have been used to recognize full plants from segmented images. Active shape models are statistical models of the shapes of objects which can be iteratively deformed to fit an object in the image. This shape modelling method has been widely used in plant classification (Sclaroff and Liu 2001; Søgaard and Heisel 2002; Johansson, Søgaard et al. 2004). The use of active shape models requires relatively high computing power. The results for weed plant recognition have so far been promising at a level of 80%. Additionally, the results of using the active shape models is highly dependent on the initial orientation of the model and the way the statistical restriction rules of shape deformation are defined for each iteration. Good results may take a lot of processing time to achieve. Apart from whole plant detection, the deformable templates or shape models have also been used in single leaf recognition (Manh, Rabatel et al. 2001). Unfortunately, single leaf recognition with shape models suffers from the same problems that whole plant shape model recognition does. Therefore, at the moment, the use of active shape models is restricted to off-line systems only.

A different approach to shape matching has been taken by Mokhtarian and Mackworth (1992) and Mokhtarian (2005). In these papers a method to represent planar curves is presented. The method of representing shapes is scale and rotation invariant but requires a full shape contour to be accessible. The representation is fairly simple to calculate and can be used to create the classification of shapes. A classification demo using their method called SQUID is freely available online (Mokhtarian 2005). This method has also been used in the application of leaf shape matching when leaves have self-intersecting parts (Mokhtarian and Abbasi 2004). Without modifications this method is not suitable for shape matching of partial shapes or shape segments, when a full shape contour is not available.

Petrakis, Diplaros et al. (2002) have developed a shape matching algorithm based on dynamic programming for partial shapes. Their matching performance is shown to be better than SQUID or other existing methods. The novelty of this method is in its ability to match partial shapes. The method is still limited such that partial shapes can only be from one shape. Therefore, the point where different objects are connected has to be known in advance. Another shape matching method using dynamic programming is presented by Gandhi (2002). This method is claimed to work with partial shapes and

overlaps. The performance of correctly recognising the shapes with these methods is at best around 80%.

Processing segmented binary images to recognize overlapping plant parts has been attempted by eroding and dilating segmented plants (Pérez, López et al. 2000). In this method the segmented blobs are eroded until only the smallest possible components, called seeds, are left and then dilated until the original size has been reached. The dilated seeds are not allowed to rejoin. The use of this method is illustrated in Figure 14.
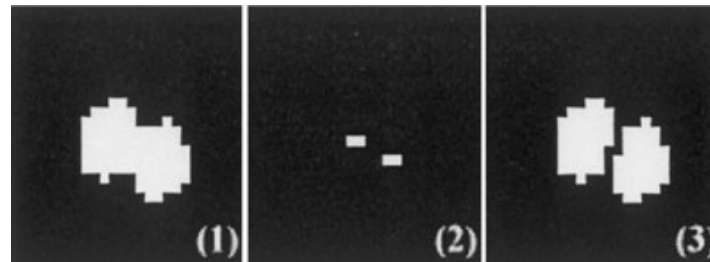
Figure 14 Example of eroding and dilating segmented blobs to find individual plant parts and leaves. (1) typical weed, (2) erosion until object division and only seeds left, (3) dilation avoiding rejoining of segments (Pérez, López et al. 2000)

Johansson and Baerveldt (2005) have conducted a comparative study on methods to overcome the problem of overlapping plant parts. A comparison of the erosion and dilation method, watershed and a new cutting method were done. Two merging approaches were also tested with the recognition or segmentation methods, and a total of six combinations of methods were evaluated.

The watershed method is based on the idea how water poured into neighbouring valleys rises and finally meets. The valleys are created by taking the distance from the edge of the blob as a topographical feature. The blob is split into segments where the water from different valleys meets. This often results in over segmented blobs.

The new method proposed by Johansson and Baerveldt (2005) is the cutting method. This process is shown in Figure 15. First all the blobs are labelled, and if a big enough blob is found it is taken as a candidate for cutting. A convex area is created around this blob (c), like a rubber band would fit around the blob, and the original blob is subtracted from the convex area. Then new areas and straight lines around the blob indicated as bases in Figure 15 (d) are created and locations in the areas furthest away from the base line are noted as in yellow circles in Figure 15 (d). Next lines are drawn normal to the base lines through the found locations. These lines segment the blob and finally the segments are merged together.
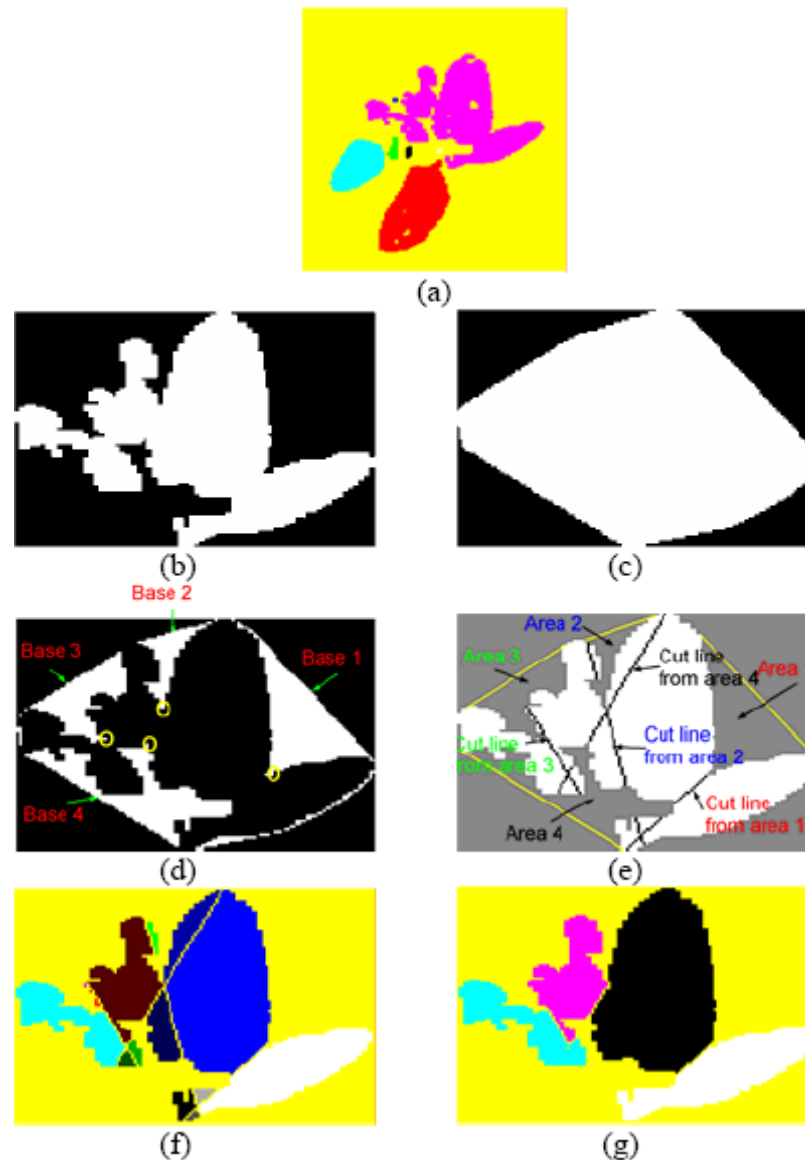
Figure 15 Cutting method to find different plant parts. (a) A labelled (in different colours) segmented image, (b) original object to be cut, (c) convex area around the original object, (d) area between leaves, (e) cut lines, (f) segments before merging, (g) segments after merging. (Johansson and Baerveldt 2005)

The cutting and merging has been found to work relatively well for segmenting overlapping parts and creating single leaves from blobs with more than one leaf in them. The watershed method has been found to produce the second best results. The cutting and merging methods were recorded to correctly mark up to 80% of the area for connected crop and weed plants.

Neto, Meyer et al. (2006) also introduced a method of finding individual leaves from overlapping leaf images. Fuzzy clustering of leaf information was used with genetic algorithms. The performance figure of extracting individual leaves was only 46%.

Overall, there is no single perfect solution to deal with overlapping plants. The

selection of the method to be used may depend on the imaging spatial resolution, plant shape and size or other such variables. It is also evident that there is a need for novel and improved methods to deal with the problem.

2.2.4 FEATURE EXTRACTION

The feature extraction step generally builds on the results from the recognition step. In some cases it is hard or even impossible to tell where recognition starts and feature extraction stops. In the feature extraction step some descriptors are calculated for the found segments. The descriptors can be derived from segment properties such as shape, colour or spectral information. Depending on the recognition method used, the features are calculated for whole plants (blobs) or individual plant parts. Examples of possible features are given in Figure 16.

Shape	**Colour**	**Spectrum**
Area Perimeter	Normalized colour	Normalized spectra
Compactness	Mean colour	Mean spectra
Elongation	Standard deviation	Standard deviation
Area		
Length/width	**Texture**	**Others**
Eccentricity	Entropy Wavelet	Plant position
Thickness	Fourier filters	Fourier shape
Moments		
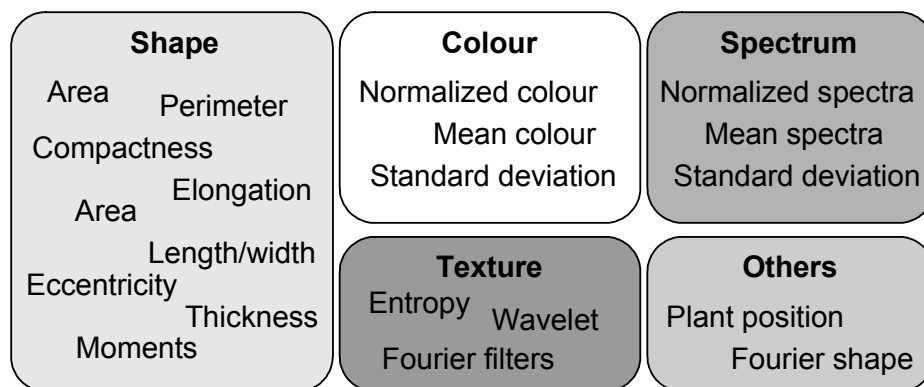
Figure 16 A list of commonly used descriptors

Varying definitions or calculation methods for similar descriptors have been used in previous research projects, as can be seen in the difference in definitions by Hemming and Rath (2001) and Åstrand and Baerveldt (2002). They both list and use a number of colour and shape descriptors. Examples of similar feature extraction processes can be also found from Lee, Slaughter et al. (1999) and Slaughter, Giles et al. (2000) in the application of tomato plant weeds. Fourier analysis can also be used to quantify leaf shapes. Elliptic Fourier analysis has been used for leaf boundary modelling from chain code by Neto, Meyer et al. (2006).

The location or pattern of the plants can give a good insight on the probability of the plant being a weed or crop. Most crop plants are grown in parallel rows with essentially equal spacing between the plants. Location and pattern information is used as a feature by Onyango and Marchant (2005), Åstrand and Baerveldt (2004) and Have, Nielsen et al. (2005).

Texture descriptors for objects represent their surface patterns. Methods such as

windowed Fourier filtering (Azencott, Wang et al. 1997) and Ridgelet transformation (S.Arivazhagan, L.Ganesan et al. 2005) are used to quantify the texture content. A comparative study of various texture classifiers was conducted by Ojala, Pietikäinen et al. (1996). A method using local binary patterns was considered to perform well and be highly efficient to calculate. Additional research with this method was done to modify the method suitable for rotation invariant texture quantification (Ojala, Pietikäinen et al. 2002).

There are no examples in literature of texture being used as a feature in single plant or leaf classification. The reason could be that leaves do not necessarily have a distinct pattern and the methods are limited to objects of a larger scale or more uniform texture. Texture classification has been only applied to plant canopy classification. For example broadleaved and grassleaved plants have been successfully classified using Gabor filters by Tang, Tian et al. (2003).

All the descriptors in shape and colour are one dimensional. Texture descriptors are also mostly one dimensional. The dimensions of spectral data descriptors depend on the used spectral resolution and spectral data range. Technically each spectral data point is its own descriptor.

Calculation times for different features vary greatly. In certain applications real-time calculations are needed and computing times need to stay within strict limits. Good features have a high descriptive value and they are fast to calculate. Determining a high descriptive power for a feature is explained in the next section.

2.2.5 CLASSIFICATION

In the final step of detection and classification, the features extracted or shapes recognized are used to estimate the plant species. Again, there are a great number of classification methods and different ways of applying them. A selection of the most commonly used methods is shown in Figure 17. The methods chosen for this project are shown in bold.

Classification ─── **Discriminant analysis** ─── **Linear**
                                                    └── Non-linear
             ├── Artificial neural networks
             ├── Fuzzy clustering
             └── etc.
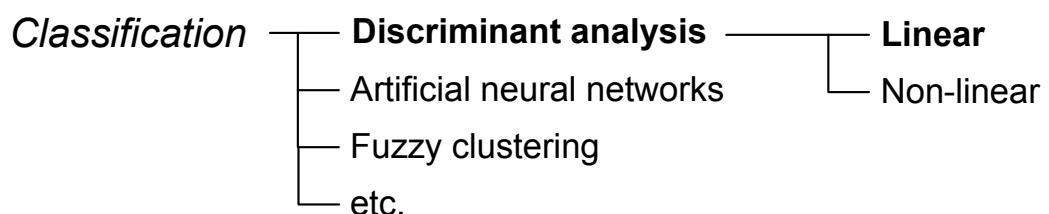
Figure 17 Classification step in weed detection and classification

The common scenario in plant classification is that prior to classification each of the plant classes are known to have certain properties or combination of feature values. When

several features are used, it is likely that some separation can be made between classes with certain combinations of these features.

Prior knowledge of plant properties can be used to create a library of plant data. One method of storing the data can be statistical, where mean and standard deviations of each feature are saved. Plant data can also consist of all known measurements of plants. The library contents, called the training set, is then used to give the classification system prior knowledge of the plants.

In the training set each plant is stored in its own class. Plant classes are discrete and non-metric since a plant can only belong to one such class. As long as the class divisions are correctly constructed, one measurement can belong to one class only.

Multiple discriminant analysis is a statistical classification method where one non-metric dependent variable is estimated from a set of independent variables. In plant classification the dependent variable is the plant class and independent variables are the features used in the classification.

The number of features used in the classification determines the dimensionality of the system. A feature set for one class typically forms a cluster of data in a multidimensional space. Classification methods then estimate the class boundaries with the given training sample and later measurements are classified according to these created boundaries.

The training set size may have a great effect on the performance of the classifier. As a rule of thumb it is suggested by Hair, Black et al. (2005) that the total number of samples in a training set should be at least 20 times the original number of independent variables. Also, the minimum number of observations per group should be 20. Groups with a larger number of observations have a higher chance of classification, and it may be desirable to randomly sample from the larger groups to compensate for the true variation in the data.

The classification example shown in Figure 18 is two dimensional, i.e. it has two features describing the training set objects. Class 1 and class 2 are fairly clearly clustered and have only some overlapping measurements. In case of a linear classifier a linear boundary between classes is created. In this case one object from each class would be misclassified from the training set. A better fit is achieved by a non-linear boundary. In the non-linear example case all the training set objects are correctly classified.

The clusters in the given example are fairly well separated, however this is not always the case with real data. More features may be needed for better separation. Also, with larger number of classes the separation often becomes more difficult.
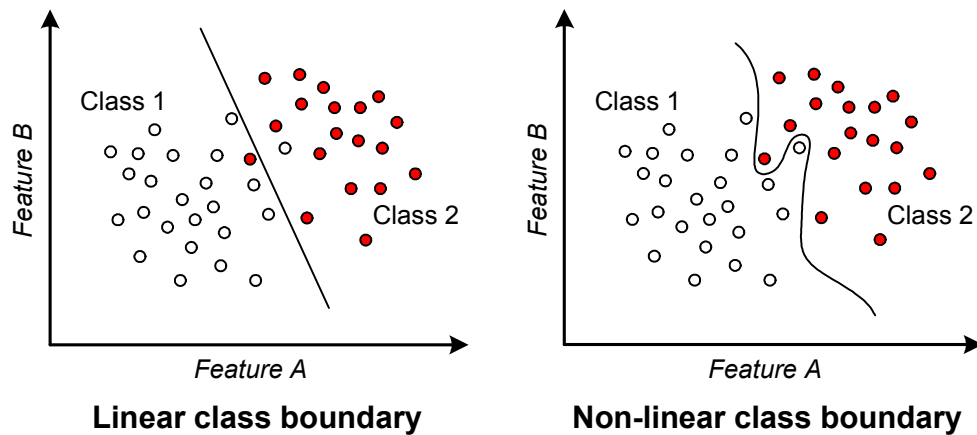
Figure 18 A two-dimensional classification example with linear and non-linear class boundaries

It is common that the decision boundary becomes highly complex when a creation of optimal boundaries between classes is attempted. A complex boundary is rarely a good generalization to the problem and only works well for the training set objects. This problem is called overfitting. In the end a simplified boundary often produces most robust classification (Duda, Hart et al. 2000).

In discriminant analysis the selection of variables can be made manually by observing the distribution and discrimination power of each variable between the classes. A beneficial variable in discriminant analysis has to separate at least one group from the others. A common method of variable selection is the stepwise selection method. Measures such as Mahalanobis distance, F-value and Wilks' Lambda can be used to estimate the discriminant power of a variable. Discriminant analysis with stepwise selection of variables can be done using software packages like SAS (SAS 2007) and SPSS (SPSS 2007). Also, a Matlab toolbox for stepwise discriminant analysis has been created by Vandev (2003).

Multicollinearity of variables should be avoided especially in variable selection of the stepwise discriminant process. Multicollinear variables are highly correlated and can be explained by each other. The addition of multicollinear variables to the discriminant functions has negative effects on the predictive ability of the functions and in the selection of further variables. The results may become biased with multicollinear variables.

Multiple discriminant analysis is probably the most common method used in applications of plant classification. Koger, Bruce et al. (2003) used discriminant analysis with low spatial resolution spectral data of approximately 0.25 m x 0.25 m per measurement. Their spectral data consisted of a mix of soil and plant spectra. Using wavelet analysis to unmix or separate the spectra and linear discriminant analysis to classify

34

them, a minimum classification probability of 87% was obtained.

Borregaard, Nielsen et al. (2000) used a line imaging spectrograph in visual and near infrared range with a spatial resolution of 1.6 mm x 6 mm. Spectral information was classified with linear or quadratic discriminant functions and crops and weeds were correctly classified with a performance of up to 90%. A similar setup with a slightly lower spatial resolution was used by Vrindts, Baerdemaeker et al. (2002) in outdoor testing under sunlight conditions. Sugar beet classification probability was recorded at 95% but maize at only 15%. A line imaging spectrograph was also used by Feyaerts and van Gool (2001) with 77-87% classification probabilities. They compared various different classification methods from nearest-neighbour classifiers to neural networks.

Neural networks are also commonly used in plant classification. Artificial neural networks are loosely modelled after the structure of the brain and consist of connected nodes or neurons working together to create an output. The networks have an input and an output layer and there may be a number of layers, called hidden layers, between them. The layers can have any number of connected neurons.

As with discriminant analysis the neural network requires training, but the methods of training are different. There are three major training paradigms: supervised learning, unsupervised learning and reinforcement learning. In supervised learning, a training set includes information about input and output pairs and can easily be used with plant training data where plant class can be known for each measurement. In unsupervised learning the output pairs for inputs are missing but the user has some knowledge on the cost functions including any combination of inputs. A model is then built to fit the observations or inputs. In plant classification, the unsupervised learning could cluster plant data measurements into different classes. In the last of the three learning methods, reinforcement learning, both the inputs and outputs are missing and are generated by interactions with the environment. An observation is created after each action and the network then changes its values according to predefined cost functions. Reinforcement learning is typically used with applications like robot control and may not be suitable for plant classification. There are several algorithms for the training of neural networks. Statistical methods and genetic algorithms are perhaps the most common. During the training process the weights of the connections between the neurons are optimized.

Three different neural network models were investigated by Burks, Shearer et al. (2005). Features for plant classification were derived from HSI colour space transformation and variables selected by stepwise discriminant process. Squares of

approximately 38 mm^2 were classified. Classification accuracies for six plant classes and soil exceeded 90% with most neural network models. The best performance of 96.7% was recorded with back propagation network topology.

Shape features were used with a back propagation artificial neural network by Cho, Lee et al. (2002). The neural network model was modified to overcome overfitting problems. A comparison of discriminant analysis and neural network approaches was conducted. Discriminant analysis correctly classified 92% of radish plants and 98% of weed plants and the equivalent neural network performance was 93.3% and 93.8%, respectively. With the method described, the overlaps between plants could not be dealt with.

A different approach to using neural networks in plant classification was taken by Aitkenhead, Dalgetty et al. (2003). They evaluated the use of self organizing neural networks for plant classification purposes without predefined features or descriptors. The performance of their system was 75%, which is somewhat lower than the desired level. However, the purpose of their research was to show the potential of a neural network approach without prior knowledge of plant descriptors.

Fuzzy classification is the third most common method used in plant classification. Hemming and Rath (2001) used weighted fuzzy classification with shape features of extracted blobs. The method correctly classified 51-95% of the plants but had problems with overlapping plants.

Meyer, Neto et al. (2004) used unsupervised fuzzy clustering to eliminate the need to manually identify the regions of interest to the system. Manual intervention is typically the most costly and time consuming issue in classification. The target was to identify plants from soil and corn and wheat residues. Four fuzzy methods were compared to manual classification. These were fuzzy excess red (FxR), fuzzy excess green (FxG), c-means clustering (CM) and Gustafson-Kessel (GK) clustering. The best results were obtained by using GK algorithms, where 16-99% of plants, soil and residue were correctly classified.

2.2.5.1 *Principal Component Analysis*

If there is a large number of independent variables and only a small number of samples the discriminant analysis will fail. The reason is that the estimated variance-covariance matrix cannot be inverted. In this case the data can be converted to principal components (PC).

Principal Component Analysis (PCA) is most often used to reduce the dimensionality of a multivariate problem (Feyaerts, Pollet et al. 1999; Moshou, Ramon et al. 2002) or more specifically to find the true dimensionality of the problem. Large numbers of input

variables are not necessarily practical for classification purposes. PCA transforms a data set to uncorrelated variables, its principal components. The eigenvalues of these components represent how much of the total variance in the data set these uncorrelated variables explain. Often the first few principal components explain nearly all of the variance in the data set. The eigenvalue order is not necessarily related to the discriminatory power of a variable and a stepwise discriminant analysis is commonly used after the PCA. (Johnson 1998)

Principal components can be calculated using Matlab or most statistical calculation tools. The corresponding eigenvalues for each principal component can be calculated from a data matrix X (variables in rows, samples in columns) by using singular value decomposition shown in the equation

$$X = USV^T,$$
(2.2)

where the diagonal values in the matrix S are the eigenvalues of the principal components. The magnitude of the eigenvalues determine the proportion of the total variance each principal component represents. When reductions in the numbers of variables are needed, it is typical to select only the principal components that represent 95% or 99% of the total variance.

2.2.5.2   *Use of Spectral Libraries for Local Plant Classification*

Spectral reflectance libraries have been created for differentiation of different objects within spectral satellite data. ASTER spectral library consists of nearly 2000 spectral samples of natural and man made objects. The library is a combination of spectra from John Hopkins Spectral Library and NASA's Jet Propulsion Laboratory measurements. (NASA 2000) The spectra in this library is collected without interest in subtle changes in vegetation spectra. These changes are not always meaningful in satellite remote sensing applications and only general spectra for grass and trees are stored. Slightly more detailed approach in spectral vegetation measurements has been taken in gathering spectral data to the USGS spectral library (Clark, Swayze et al. 2007). In this collection of spectra, a number of plant species have their own library record but the spectra is collected always from only one sample in laboratory.

If a meaningful library would be constructed, the spectral variability due to different growing conditions (irrigation levels, available nutrition, soil etc.) and growth stages needs to be included. Plant spectra between different species are really similar and even the smallest changes in growing conditions can produce significant changes for classification

purposes. (Lewis 2001; Fyfe 2004)

Due to several variables affecting the plant spectra, it is nearly impossible to create a reliable library for all possible plant species and their spectral variations. It is perhaps possible to create spectral libraries for local conditions and a limited set of plants. This approach would include measuring the spectra of plants at a certain time at a certain location. Reliable use of this created library would then happen locally within a limited timeframe.

2.2.6 DISCUSSION ON CURRENT PLANT CLASSIFICATION RESEARCH

Previous research projects have shown that no single detection or classification method is superior to another. However, certain features in the systems can be thought more beneficial than others. The ability to detect the red edge created by the chlorophyll reflectance change in the near infrared region is a simple indicator of plant material and its detection can be achieved with simple filters. Narrowband spectral data can further assist in distinguishing between different plant classes, as they all have unique spectra depending on the growth stage and conditions.

Addition of the knowledge of shape to a classifier is shown to be valuable, but it is difficult to measure in real growing conditions where plants often overlap with each other. These overlaps create problems for systems attempting to fit complete shapes around segmented blobs. If a system would be reliable in separating the overlapping parts, a partial shape classification could be utilised. No tool to make the separation is currently available, and therefore the implementation of this proposal would require further research.

The accuracy classification with current systems is good enough for reductions in chemical use compared to broadband spraying. Then again, systems capable of robust classification at leaf level under varying conditions are not available. Systems may work well with certain plant combinations, in previously known plant patterns or lighting conditions. Problems often arise with overlapping plants and in changing light. Also, most current systems are optimised and built for a particular crop-weeds combination and are not necessarily flexible enough to work with any plant combination. Flexible detection and classification system would be able to correctly classify all plants separately, although in many cases the 2-class classification to weeds and crop is sufficient. Separation of different weed types would allow specific targeting of weeds in the methods most harmful for them. The current method of spraying is to mix suitable chemicals together and spray all weeds with this mixture. Savings could be achieved by matching the right chemicals with the right weeds.

2.3 REQUIREMENTS FOR AUTONOMOUS WEED DETECTION AT THE SINGLE LEAF LEVEL

A number of automatic weed detection systems have been developed, but only the most simple systems such as WeedSeeker and GreenSeeker from the Californian company NTech Industries (NTech Industries 2003) or Detectspray originally from Australia have been commercially released. These systems rely on the low resolution reflectance of infrared or a combination of infrared and visible light and are only capable of detecting green vegetation. Individual plant classification is not possible with these devices. However, savings in herbicide costs up to 80% have been recorded with these products compared to uniform spraying. When tested in varying conditions, WeedSeeker performed the best with high density broadleaved weeds (Antuniassi, Nery et al. 2003).

For greater flexibility, accuracy and further reduction in chemicals or the ability to precisely locate weed plants, more advanced systems are needed. A good classification detection system needs to be able to separate weed and crop plants at the single plant or leaf level. This leads to a specification that the system needs to recognize any combination of plants with reasonable accuracy.

Autonomous weed detection is the only viable solution for automated weed control. Manual weed detection is not feasible and even then 65-85% of the weeds can be left unnoticed (Slaughter, Giles et al. 2000). Lee, Slaughter et al. (1996) showed that the cost of spraying and hand weeding was $50 and $80/acre (0.4 ha), respectively. Economic analysis and a prototype machine showed that the purchase of a $110,000 machine would be justified if this robotic system could travel at 0.80 km/h (0.5 mi/h). This calculation assumes a three-row machine with row spacing of 1.52 m (60 inches) and an operating period of 45 days per season, 60% of overhead and operating costs, no interest, and a five year machine life.

Onyango, Marchant et al. (2005) conducted simulations using crop and weed competition models indicating the optimal levels of correct classification for crop and weeds to allow certain crop yields. Their results show surprisingly that minimising the crop damage (misclassified crop) or maximising the weed removal (correctly classified weeds) does not automatically result in the best yield. It is suggested that an optimal trade-off between the classification probabilities can be found. In their case the optimal weed detection probability for maximum yield was 80% and approximately 93% for crop detection.

Based on the information gained from previous research projects, the following points can be thought of as minimal specifications for an autonomous weed detection system:

1. Accurate
   - Spatially accurate at sub-centimetre level
   - Accurate detection and classification of plants types at >80%
2. Flexible
   - Adaptable to any plant types
3. Low-cost
   - Purchase and operating costs cannot exceed current price of weed control
4. Quick
   - Real-time operation desirable with immediate control action after detection

Single leaf level classification will require an increase in spectral imaging spatial resolution compared to the previously built systems. Also, the operating speeds need to be increased in order to achieve economical feasibility. The presented points should be considered when new techniques for plant detection and classification are investigated.

2.4 RELATED SYSTEMS AND APPLICATIONS

This section gives examples of systems related to the application of weed classification. A complete listing of such examples would be endless and only the examples thought most relevant for an overall picture are given. When an autonomous device for weed classification is considered, satellite based positioning and crop row following can be considered as supporting systems. Their contribution is separate from the actual classification and control, but is often vital for the navigation tasks. Sensor fusion in remote sensing and material sorting examples are given as they represent a classification system with similarities to a weed classification system.

2.4.1 SATELLITE BASED POSITIONING

Satellite based positioning, generally based on the GPS system, is commonly used in localisation of robots, vehicles, airplanes, etc. The system is reliable in any outdoor conditions and has a predictable base accuracy of 5 m to 10 m (Ashkenazi, Park et al. 2000). There are also plans for the launch of a similar European controlled system called Galileo in 2013. Its accuracy will be under 1 m. Accuracies of both systems can be improved by using a ground beacon at a precisely known location. These are called DPGS or RTK-GPS (Real-Time Kinematics) systems. With a ground beacon, accuracies of just few centimetres can be achieved.

Gerhards, Sökefeld et al. (2002) used the accurate DGPS system to spray weeds with the aid of premade weed maps. Griepentrog, Nørremark et al. (2005) have mapped sugar beet seed locations with an accuracy of 16-43 mm. The seed locations were later used as an input to an optical plant recognition system improving the recognition probabilities. Have,

Nielsen et al. (2005) have also used a GPS based system to map Christmas tree locations and guide an autonomous vehicle mowing around the tree trunks. These examples in agricultural applications show the variety and potential available in satellite based positioning.

2.4.2 CROP ROW FOLLOWING

Plenty of research effort has been put into the automation of following crop rows. The subject is important for drivers of farming equipment and for the quality of work. It requires high levels of concentration to steer a tractor or a harvester in a straight line following the crop rows. In weed classification and control the crop row following could be used in steering an autonomous vehicle equipped with the detection system along the crop rows.

Billingsley and Schoenfisch (1995) proved a simple prototype system based on one monochrome camera could achieve 20 mm row tracking performance. The same accuracy was obtained by Marchant , Hague et al. (1997) using a single monochrome camera in NIR to find crop rows. This performance was determined to be satisfactory accuracy for autonomous guidance. Åstrand and Baerveldt (2005) have developed a similar crop row recognition system based on Hough transform to find crop rows and a steering wheel to guide the cultivator. They recorded slightly better accuracies of between 0.6 cm and 1.2 cm and verified the method with both a tractor and an autonomous vehicle. Accuracies within similar ranges are recorded with other systems with varying image processing techniques (Tillett, Hague et al. 2002; Søgaard and Olsen 2003; Han, Zhang et al. 2004). The completed projects show that the accuracies in row detection and following are at a desirable level and that they can be achieved in several different ways.

Stereo cameras have also been investigated in crop-row detection (Kise, Zhang et al. 2005). The benefit of such systems is the added range dimension. If single camera solutions fail for example in the case of heavy weed densities between the rows, the performance of a stereo system may not be affected.

Current row following accuracies generally exceed the satellite positioning accuracies. A combination of a global position from the satellite systems and a precise relative position and guidance from a row following system increases redundancy and reliability.

2.4.3 HYPERSPECTRAL AND HIGH RESOLUTION COLOUR FUSION IN REMOTE SENSING

A simulation of fusing hyperspectral and relatively spatially more accurate colour data is presented by Hsu and Burke (2003). They give examples of current methods in registration

and spectral sharpening using the data from both sensors.

The conclusion of their paper is that a combination of hyperspectral data and higher resolution colour data can be used to improve classification compared to using only lower spatial resolution hyperspectral data. The accurate registration of data, i.e. knowledge of how to locate the same area or pixel in both spectral and colour data, is the key for improved results. Colour data provides accurate border segmentation of the objects while hyperspectral data provides discriminating information for robust classification of different materials.

This combination of spectral and colour imaging could be adapted for use in plant classification. A colour camera could provide accurate information on plant borders and their locations and spectral data further enable classification to different plant types. Although the combination of colour and spectral imaging in remote sensing is common, the direct application of the analysis methods in the macro imaging (spatial resolution of a metre or some metres) to images with resolutions of few millimetres may not be straight forward. The reflectance properties of a region in a single leaf compared to canopy of plants is considered fundamentally different.

2.4.4 MATERIAL SORTING

The Austrian research institute CTR has developed a commercial product for material sorting based on spectral imaging (Kulcke and Kerschhaggl 2006; CTR 2007). The system is capable of recognising different plastic materials on a conveyor belt travelling at speeds of up to 2 m/s and sorting the materials in different bins with the help of an array of pneumatic valves. The recognition is based on spectral reflectance measurements. The online detection rate at high speeds is based on high frame rates of up to 200 fps and real time calculations using field programmable gate array (FPGA) hardware.

Although the classification of different plastic materials is easier than classifying plant spectra, the example of material sorting could be adapted to classification of plants. Conveyor belt solution is obviously not possible on fields and cameras would need to be moved instead.

2.5 SUMMARY OF WEED CONTROL RESEARCH

An overview of the current weed control methods and technologies involved in automatic weed detection has been given in this chapter. Special attention has been drawn to the imaging systems and data processing methods to extract weed locations in fields. The limitations in previously built optical detection systems have been outlined and a set of

requirement suggestions for future systems were given.

Automatic mechanical weed control methods are limited to treating between-row weeds without accurate within-row crop and weed plant detection systems. In chemical control the lack of within-row plant detector results in unnecessarily high volumes of herbicides being sprayed in fields. Economical and environmental savings could be achieved if a reliable system locating the weed plants would be built.

3.0 PROPOSED PLANT DETECTION AND CLASSIFICATION METHOD

Several studies have shown that the automatic plant detection and classification with cameras is feasible, although as explained in §2.2.6 there are still unsolved problems. Most researchers have used either colour or spectral systems for plant detection.

Slaughter, Lanini et al. (2004) have conducted an evaluation of colour and spectral classification of tomato plants and weeds. Discriminant analysis was used in classification, and five spectral bands in the range of 400 – 2320 nm were used. Spectral imaging in the range of 400 – 1000 nm was done with a silicon sensor. Sensors for wavelengths above that are specialized and more expensive. The results from the study suggest that spectral imaging at most wavelength ranges outperforms the conventional RGB imaging. This conclusion is in line with many other research projects as well.

Another and perhaps more relevant issue investigated in Slaughter, Lanini et al. (2004) was the effect of image bit depth to the classification accuracy. Data with 16, 14, 12, 10 and 8 bits were compared in the study. Obviously the results were better with a higher bit depth, and the difference between 8- and 16-bit data was significant but only for spectral data. There was also a big difference between 10- and 8-bit data. Cameras with over 10-bit depth are relatively expensive and judging by the results of the study, the 10-bit cameras may be the solution to optimize accuracy and price. Cameras with 10-bit analogue to digital conversion and 8-bit image data were chosen to be used in the presented project without compromising the data accuracy for cost.

The proposed plant detection system consists of two separate imaging devices with simultaneous image acquisition. RGB colour camera and a monochrome camera with the chosen spectrograph both image at a high spatial resolution having the same objects visible to both cameras at all times. This novel combination of a spatially accurate spectrometer and a conventional colour camera has the potential to perform robust and accurate plant detection and classification. Brown, Jackson et al. (2000) have used a comparable setup in an application of printing quality analysis. Although the idea of combining these two imaging technologies is not novel, the similar camera setup with a high spatial accuracy has not been previously used in agricultural applications.

The spatial accuracy of the colour camera in the proposed setup is still approximately five times more accurate than the spectrometer spatial accuracy. Fusing the information from these two imaging devices using their individual benefits is shown to improve the overall results compared to using either imaging device individually (Hsu and Burke 2003).

The chosen spectrograph was a Specim Imspector V10 with a nominal spectral range of 400 – 1000 nm. This spectrograph was a line scan type spectrograph imaging a narrow line and dispersing the spectral components in that line to the vertical axis of the image. This technology is comparatively low cost, compact and produces accurate information in both spectral and spatial domains. This spectrograph type has also been chosen by other researchers in plant detection (Borregaard, Nielsen et al. 2000; Vrindts, Baerdemaeker et al. 2002).

Colour and spectral imaging technologies have already been individually proven to be capable of providing satisfactory plant detection and classification results under certain conditions. Therefore, it is likely that the results will be improved with the two technologies combined. This allows the best possible feature combinations to be chosen for classification purposes and the system becomes flexible such that one system can be used with any combination of plants with minimal changes.

It can be argued that the same information can be gathered using tunable filters but at the moment the filter technology is much more expensive than using two cameras and a line imaging spectrometer. Tunable filters have also other disadvantages for use in applications involving movement. Since the combined data from colour and spectral cameras has not been utilized before, this creates opportunities to create new and powerful algorithms for plant detection.

Two imaging devices in the detection system obviously produce more data than a single one would. Acquiring two high resolution images simultaneously heavily occupies the data transfer medium and increases the processing times. Without proper care this flood of data may make real-time processing an impossible task. However, it should be noted that acquiring two 8-bit images produces as much data as a single 16-bit image would.

The basic concept of the proposed novel dual camera detection and classification system is to acquire simultaneous pairs of images from the colour and spectral cameras. The selected line imaging spectrograph needs to be moved spatially to gather the full spectral cube of information. The cameras are fixed in position relative to each other and moved gradually along in the direction of the vertical colour image axis (colour image y-axis in Figure 37). Image pairs are acquired at a set interval from a predetermined and unchanging distance to the ground and colour images are used to estimate the movement by spatial registration of consecutive image segments.

A flow chart of the main image processing tasks involved with the proposed two camera solution is shown in Figure 19. The individual processes involved in data

collection, storage and processing will be explained thoroughly in this chapter and the figure is referred to where applicable. Gray areas at the top of the flow chart indicate the separate cameras and the possibilities to process the data separately. This could ultimately allow part of the processing to be done inside the cameras so that only the necessary information will be transferred via the interface link to the central processing unit. This is one example of avoiding the problem of having too much data to process.
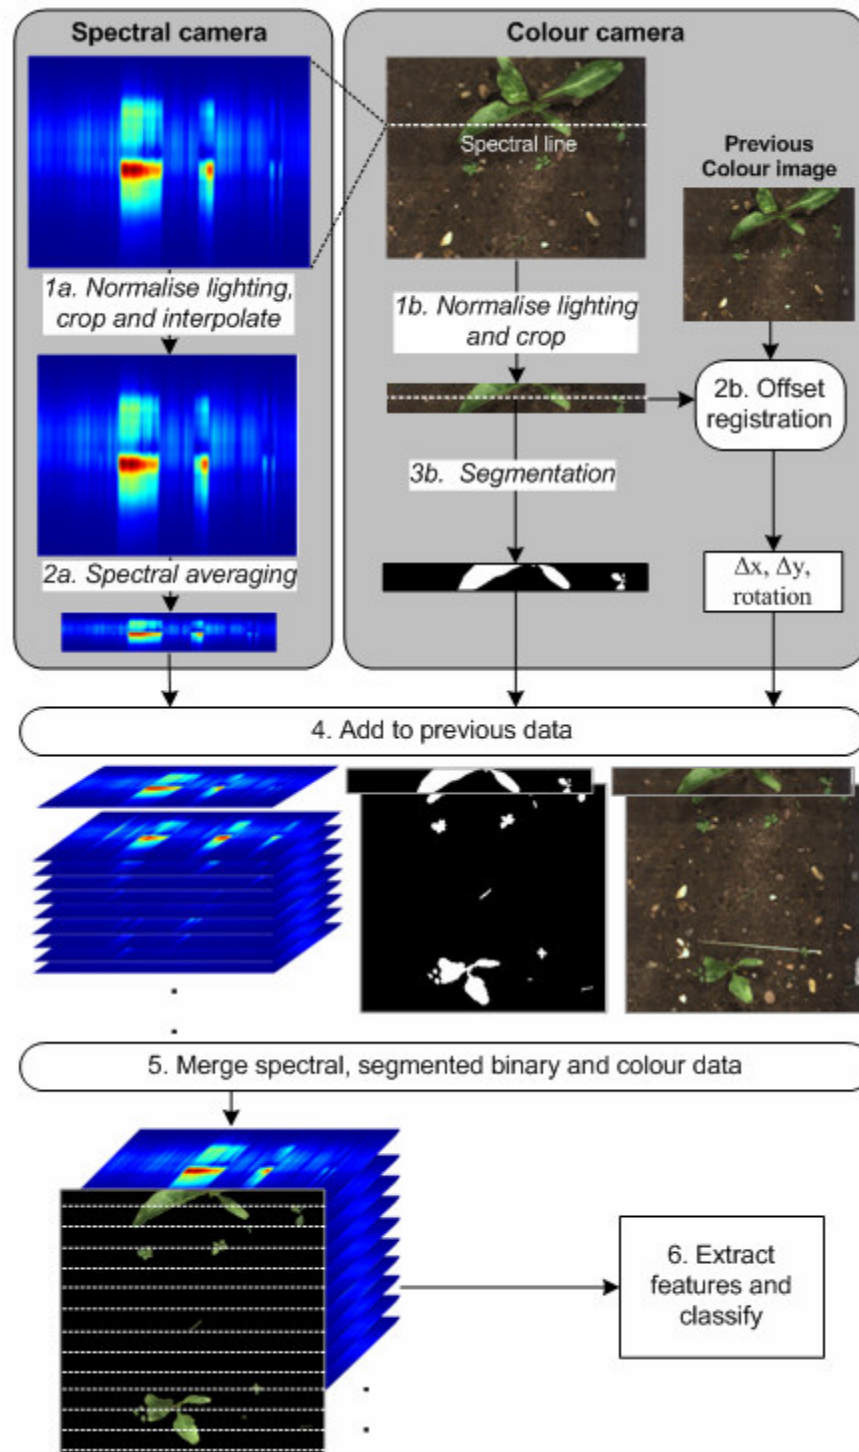
Figure 19 Image and data processing flow chart

Two different data processing approaches for creating the classification data are also discussed. These are full leaf method and window method. In the full leaf method the classification is based on colour, shape and spectrum characteristics of found full leaves. The novel window method does not consider shape properties but attempts the classification of small image areas (windows) with colour and spectral data only. The window method is also an attempt to overcome problems with connected or overlapped plant parts.

Process 6. in Figure 19 is perhaps the most important of all the processes. This is where the pre-processed collected data is fed into the classification algorithm. The proposed classification method is discussed later in this chapter.

3.1 COLOUR DATA PROCESSING

3.1.1 LIGHTING NORMALISATION IN COLOUR IMAGES

The lighting intensity is rarely uniform throughout the whole image area. In applications where accurate colour measurements or sensitive thresholding procedures are done it is vital to calibrate the system to these intensity changes. Knowledge of the intensity changes can be gathered in calibration images. These calibration images are taken of a uniform material with constant reflectance properties. Normalisation of the acquired images can be done by dividing the pixel intensities with the corresponding calibration image pixel intensities. The process of creating these calibration images and using them is further explained in §5.1.6.

3.1.2 COLOUR IMAGE REGISTRATION

The test setup was designed such that no data of the actual global location of the camera system is recorded, although the stepper motor always moved the cameras by a same distance per step along a linear guide. Image acquisition happened after certain number of steps and most of the colour image was always overlapping with the previously captured image. The movement of the cameras relative to the ground can therefore be estimated from the images. Since the camera system behaves like a scanning device moving along the vertical colour image or y-axis, a continuous image could be created from a set of overlapping images.

The process of colour image offset registration (2b in Figure 19) estimated the offset of the latest and previously acquired images within a pixel by means of cross correlation. A single line from the new image was compared to the lines in the previous image and a maximum correlation of a line pair indicated the offset required for registration. The offset

was always in the y-direction only, as this was the only direction the platform was moving. Similar cross correlation method could be modified to give x-, y- and rotational offset values. The new image was overlaid on the previous image using the estimated y-axis offset to create a continuous image (4. in Figure 19). Only small vertical segments of the red channels in the latest and previous images were used in the cross correlation process. This was done to speed up the process compared to the three-channel cross correlation.

3.1.3 PLANT SEGMENTATION

After colour image registration, the plant material from the newest portion of the image needs to be found. This process is indicated as 3b in Figure 19.

A test of six different segmentation methods was done to find out the best colour space transformation or combination of transformations for plant segmentation for the equipment used. These methods included RGB, HSI, IV_1V_2, $i1i2i3_{new}$, RGB+HSI and RGB+$i1i2i3_{new}$ transformations. The first four methods used three image channels each and the last two used six image channels for segmentation. HSI transformation was calculated with a Matlab function `rgb2hsv.m` and IV_1V_2 and $i1i2i3_{new}$ with the following equations:

$$\begin{bmatrix} V_1 \\ V_2 \\ I \end{bmatrix} = \begin{bmatrix} -\dfrac{1}{\sqrt{2}} & \dfrac{1}{\sqrt{2}} & 0 \\ -\dfrac{1}{\sqrt{6}} & -\dfrac{1}{\sqrt{6}} & \dfrac{2}{\sqrt{6}} \\ \dfrac{1}{\sqrt{3}} & \dfrac{1}{\sqrt{3}} & \dfrac{1}{\sqrt{3}} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{3.1}$$

$$\begin{bmatrix} i1_{new} \\ i2_{new} \\ i3_{new} \end{bmatrix} = \begin{bmatrix} 0.34 & 0.33 & 0.33 \\ 0.07 & 0.39 & -0.54 \\ -0.35 & 0.51 & -0.14 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{3.2}$$

A test image of different plant leaves and other objects was taken under controlled and unchanging lighting conditions. A binary reference image was created from the test image by manually marking the leaf material as one and everything else as zero. The original test image and the reference image are shown in Figure 20.

(a)   (b)

Figure 20 (a) A test image and (b) a reference image for colour transformation evaluation in plant segmentation

For each channel, the boundary values or threshold limits were manually set observing pixel values at leaf locations. A low and a high threshold values were set for each individual channel such that after excluding pixels below and above the low and high threshold limits, respectively, none or minimal plant material was excluded in each channel. Logical AND-operation was then used to combine the channels. The process of thresholding individual channels and combining their results with an AND-operation is illustrated in Figure 21. In the given example, the hue channel separates the plant material and intensity channel the bright non-plant material best.

The presented thresholding method was adopted to minimise the loss of valuable plant pixels during the segmentation process and the manually set limits allowed false positives. If a non-plant pixel was not present in one or more of the individual channels after thresholding, the logical AND-operation excluded this pixel from the combined result.
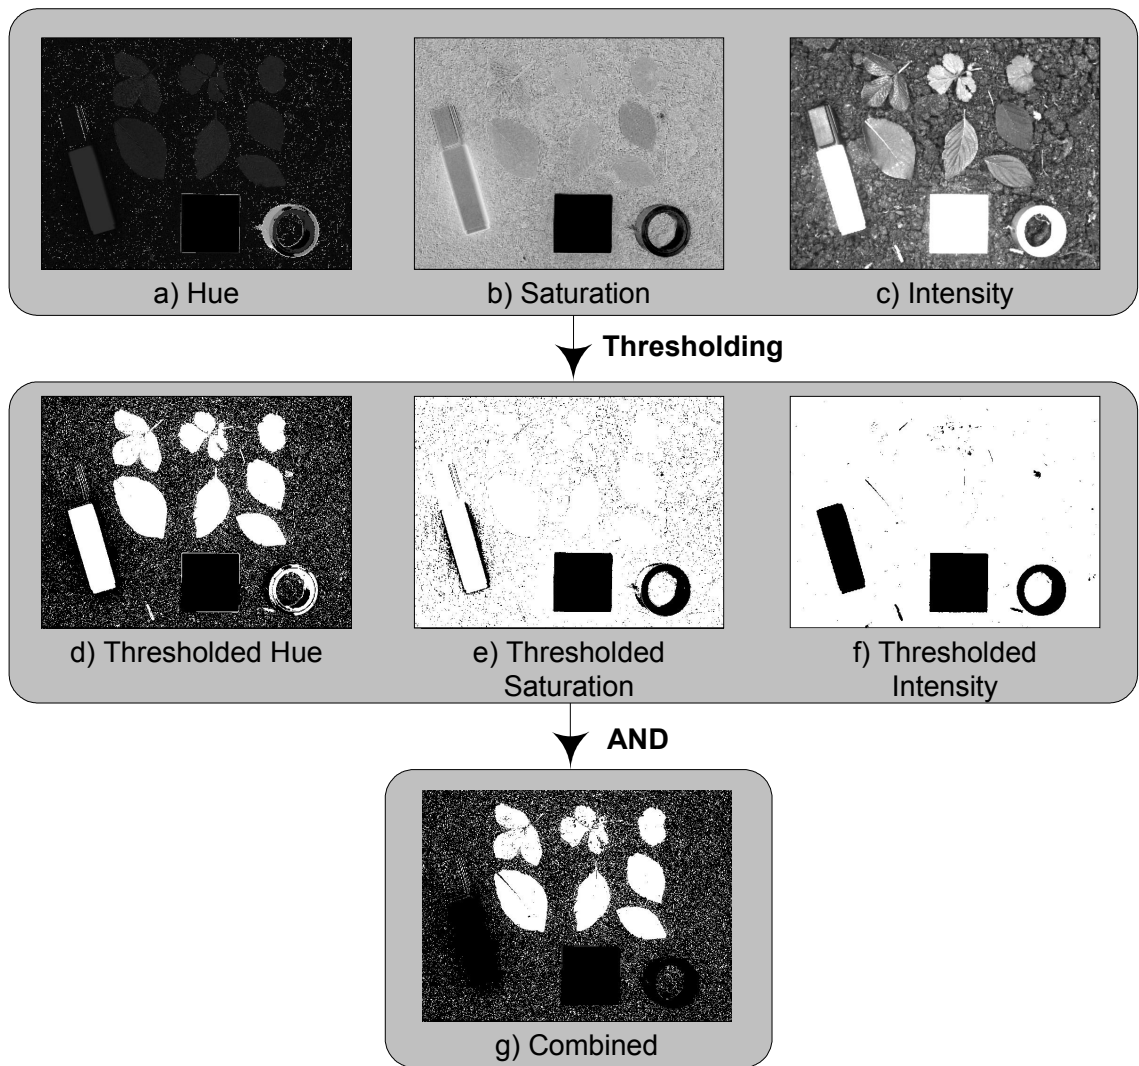
Figure 21 A segmentation example showing separate channels of HSI transformation and the resulting image after combining thresholded channels with an AND-operation

Combined thresholding results, as in Figure 21 g), using the presented six transformation channel combinations are shown in Figure 22. RGB results in a) seem considerably worse than with other transformation combinations. When threshold limits were set for RGB channels, the soil could not be reliably excluded in any of the channels. The green channel contributes most to the overall result, but red and blue channels fail to improve the result. It seems that the colour vectors in the RGB cube are not pointing in the right direction in terms of segmenting the plant pixels with the presented method. More sophisticated ways to define plant data boundaries could be beneficial for use with RGB-space. However, it seems that other transformation combinations illustrated in Figure 22 b)-f) give relatively good results.

(a) RGB   (b) HSI

(c) IV_1V_2 (d) $i1i2i3_{new}$

(e) RGB+HSI   (f) RGB+$i1i2i3_{new}$
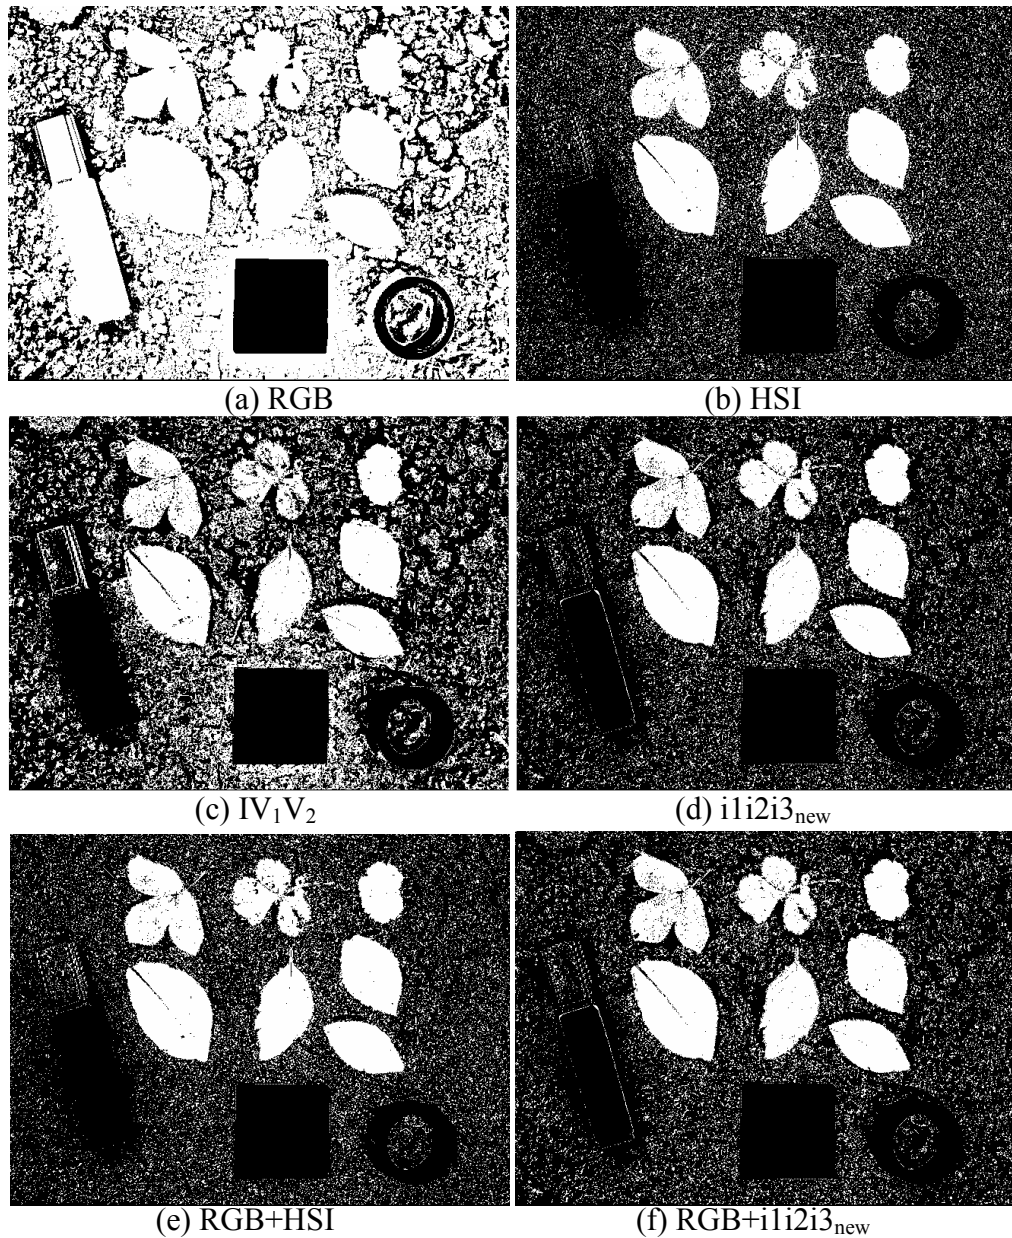
Figure 22 Segmentation results of six different methods

The results after segmentation were noisy and some post processing was needed to extract only the green plants from images. The post-processing of segmented images is illustrated in Figure 23 using the test image from Figure 20 as a source. First, the binary images were filtered with a median filter that effectively reduced the salt and pepper type of noise in the segmented images. After median filtering each output pixel contains the median of a selected pixel neighbourhood. After careful exploration of various filter sizes, a 4 by 4 neighbourhood was found to work best for the images. The resulting filtered image contained values of 0.5 and was converted back to a binary image, discarding all but the pixels with a value of one.

51
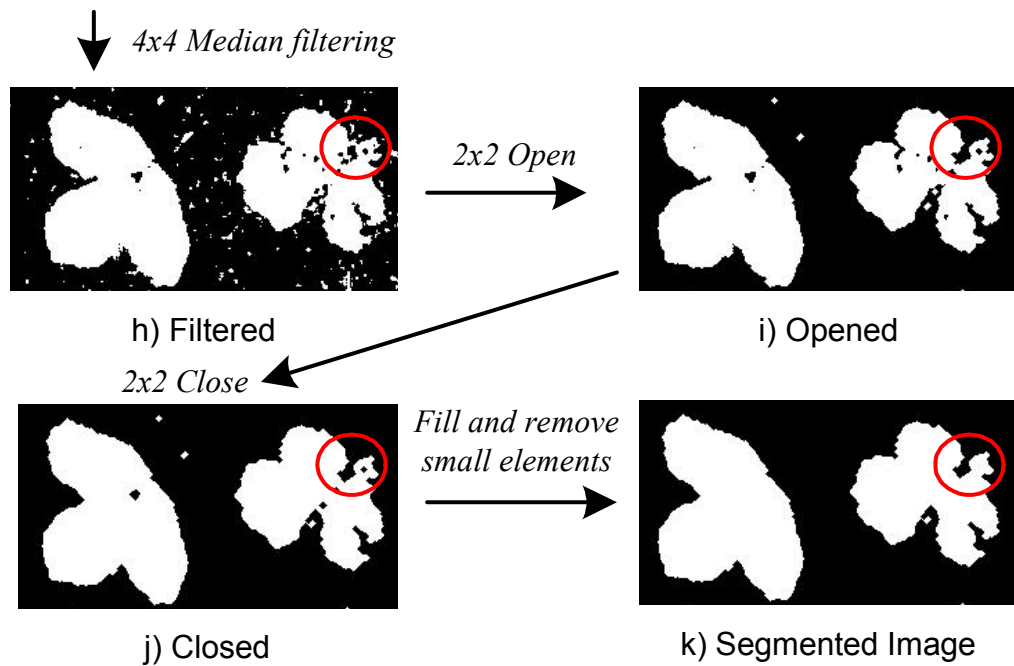
Figure 23 Post-processing example of segmented images

Some noise was still present in the images after median filtering. First a morphological open (erode followed by dilate) and then close (dilate followed by erode) operations with 2 by 2 pixel elements were performed to further improve the match with the reference image. An open operation removed excess pixels especially in the leaf borders and opened up small gaps naturally present in some leaves. Performing a close operation after the open operation with the same size element attempted to restore the original pixels before opening. After open and close operations, all the remaining small holes in the image were filled and elements or blobs smaller than 100 pixels were removed. An area indicated by a red circles in Figure 23 shows how the natural gap in the leaf is restored using open and close operations.

The method of finding the suitable processing algorithm or selection of image processing tools was done by exploring various filtering and morphological imaging options, while visually or computationally observing their performance compared to the ground truth data shown in Figure 20 b). Although the process produced good results, they are not necessarily the optimum in operational accuracy or computational load.

A performance figure for each segmentation method was calculated by comparing the resulting segmentation to the reference image pixel by pixel. The calculation compared the number of correctly segmented pixels in the whole image to the total number of pixels. Performance values for methods with and without post processing are shown in Table 1.

Table 1 Performance values for six segmentation methods

Method	Performance before post processing (%)	Performance after post processing (%)
RGB	44.9	26.1
HSI	87.2	99.1
IV_1V_2	79.1	94.3
$i1i2i3_{new}$	89.8	99.3
RGB + HSI	89.1	99.3
RGB + $i1i2i3_{new}$	89.8	99.3

Apart from RGB and IV_1V_2, the methods performed almost perfectly. No significant difference can be seen with the remaining four methods. Using RGB with $i1i2i3_{new}$ did not improve the segmentation results for that method.

The method of using RGB and HSI channels was selected. Since images were acquired in RGB format, no additional calculations were needed to use these channels. Also, Matlab has a fast function for RGB to HSI conversion. Segmentation with $i1i2i3_{new}$ channels could be used in a situation where this method would be faster to calculate. An example of segmentation results with the RGB+HSI method after post processing is shown in Figure 24. The red areas in the figure represent misclassified pixels and are mainly concentrated around the edges of the leaves. Misclassification around the edges is partially due to subjective manual selection of reference areas. No pixels were misclassified outside of the leaves.

Figure 24 A segmentation example with RGB+HSI after post processing. Misclassified pixels in red

None of the segmentation methods using colour images only can detect the difference between living green plants and other green objects. However, spectral data can later on be used to differentiate between living plants and other green material potentially segmented as plants.

3.1.4 COLOUR IMAGE FEATURES FOR CLASSIFICATION

A number of features or descriptors were designed to give numerical representation to colour and shape properties of segmented leaves or plant parts. Shape descriptors were calculated from segmented binary images and colour descriptors from colour images in the segmented regions only. The descriptors were calculated for classification purposes. Several different descriptors were created and tested so that the best ones could be selected into the final classification algorithm.

Descriptors were calculated for each separate area, called a blob, in the segmented image. Blobs were detected only when they were fully visible. Colour descriptors included six values calculated from the red, green and blue channels. The channels were first normalized dividing the channel value with the sum of all three channel values at each pixel location. Mean value and standard deviation descriptors for all three channels were calculated using these normalized pixel values.

The shape descriptors used include area, perimeter, thickness, elongation and compactness as in Åstrand and Baerveldt (2002). In addition, eccentricity was also calculated. The area value is a simple count of pixels in a blob and perimeter is the number of pixels on the boundary of the blob. Thickness is defined by the number of shrinking steps it takes to reduce the blob until no pixels remain. Eccentricity, a measure of ratio of minor axis to major axis of a bounding ellipse, has a value from zero to one and it represents how round or eccentric the blob is. An ellipse with the resulting eccentricity value has the same geometric second moment as the blob and was calculated with a Matlab function `regionprops`. An eccentricity value of zero means that the blob is a perfect circle and a value of one represents a line segment. Elongation represents a similar measure but is defined as the area divided by thickness squared. Compactness is defined by area divided by perimeter squared.

The calculation of shape descriptors does not take into account or detect any overlapping or connected plant parts. Prior to the calculation of shape descriptors it is assumed that all the blobs represent a full leaf.

3.2 SPECTRAL DATA PROCESSING

Spectral images are monochrome 8-bit images with 1360x1024 pixel resolution. The horizontal axis in the image holds the spatial data and the vertical axis the spectral data. In theory the spectral range of $400 - 1000$ nm is distributed evenly along the vertical image axis resulting in spectral resolution of 0.59 nm/pixel at the image sensor. A 4 nm spectral resolution was given in the spectrograph specifications.

The spectral image spatial axis was cropped and resized to match the colour image columns. This is labelled as 1a in Figure 19. The amount of cropping and resizing was constant and was determined by a calibration process where a colour image and a spectral image were registered together. This calibration procedure is explained later in §5.1.4.

Spectral images were also resized in the spectral axis to match the colour image height of 493 pixels, resulting in a theoretical spectral resolution of 1.22 nm/pixel. This resizing was done mainly to reduce the amount of data without sacrificing the spectral resolution accuracy. The spectral data in the resized image was still over three times more accurate than the spectrograph manufacturer's spectral accuracy specifications.

Spectral image columns represented the spectral measurement at a particular location. Before processing the cropped and resized image further, the image was adjusted according to the current lighting conditions. During normalization the data in the spectral image columns was divided by values obtained from lighting intensity measurements. This process of lighting calibration is explained thoroughly in §5.1.6.

After the lighting adjustment the spectral image columns are individually normalised by subtracting the mean of the column from each data point and dividing the values by the standard deviation of the column. This method of normalisation results in relative measurements. These relative measurements are robust towards expected intensity changes resulting from plant leaves being imaged and illuminated at different angles.

Finally, the number of spectral points was reduced by calculating average values over n pixels. Changing the value of n effectively changes the spectral measurement band width. Several different values of n were used to find the best one for classification purposes. Data averaging is indicated by 2a in Figure 19.

3.3 OVERALL DATA COLLECTION AND MERGING

The process of gathering data happens an image pair at a time. The colour and spectral images were acquired simultaneously and transferred to the computer. There was a constant spatial offset between the image pairs and this offset was determined during the colour image registration process.

The new portion of the colour image was added to the top of the old colour images, creating a continuous "tall" image. The new colour image portion was also segmented and added to the top of the "tall" segmented image. Spectral data location in the colour image was registered for each image pair and previous locations updated accordingly. These spectral data line locations are indicated in Figure 19 by dotted white lines.

Overall, the previous colour image, the continuous colour image, the continuous

segmented image and the spectral data gathered were stored in the memory. The data in the continuous images and spectral measurements were used to find segmented blobs for classification. These blobs were recorded only when they were fully visible, i.e. they were no more connected to the top line of the continuous image. The amount of memory used to store the images was controlled by setting a maximum height for the continuous images. The height limit effectively determined how many image pairs were kept in the memory.

3.3.1 WHOLE LEAF METHOD

The first of the two detection and data processing methods examined was called the whole leaf method. In this method it was assumed that one blob represents one leaf and no overlapping leaves were present. The method was expected to work well in laboratory conditions where leaves could be placed separately under the cameras. In this method colour, shape and spectral data was calculated and stored for each detected blob.

Spectral data for blobs was stored only from locations where the data is non-mixed. This non-mixed restriction means that the spatial area covered by a spectral data point had to be completely inside the blob in order to be registered. Each spectral data point was rectangular in shape and its size was governed by the spectrograph spatial resolution.

Non-mixed data was ensured by eroding the blob in the horizontal and vertical directions by the known spectral data point dimensions and storing only the spectral data points whose centre point falls inside this eroded blob. In the whole leaf method the spectral data was averaged over the whole leaf.

An example colour image set created from 20 images is shown in Figure 25.a. Data from the first image is at the bottom of the image. Figure 25.b shows the segmentation results from this colour image such that only blobs that are fully visible are segmented. Horizontal black lines in the segmented image show the spectral line locations and the number on the segmented leaves indicates the order in which the leaves have been detected. Colour and shape descriptors and spectral data have been calculated for each of the blobs shown in the segmented image.
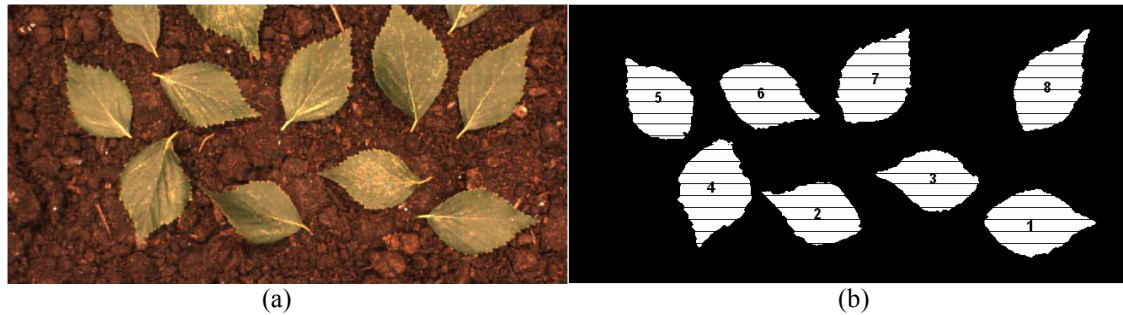
Figure 25 a) A colour image example from 20 image pairs and b) its segmented equivalent. Spectral data line locations shown as horizontal black lines.

3.3.2 WINDOW METHOD

Since detection of separate leaves cannot be guaranteed in all cases in real field environments with living plants, a method to manage overlapping leaf parts was created. The window method was a novel approach and an attempt to try and avoid problems that occur when leaves are overlapped or only partially visible. The spatially accurate optical system allows data processing in sub-centimetre level.

The basic idea of the window method is to divide the images into small windows for which spectral data and colour descriptors are calculated individually. Shape descriptors are not used with this method. The size of the windows is determined primarily according to the expected distance between acquired images and the spectrograph spatial resolution. The windows are defined as squares. In essence, the method down samples the colour image spatial resolution and spectral image spatial axis data to achieve a nearly square grid of data in each image pair.

The method of extracting windowed data itself is not novel within plant detection. For example Burks, Shearer et al. (2005) have classified colour data windows of 38 mm^2 as explained in §2.2.5. The novelty of the approach is the accurate spatial resolution of the windows and the inclusion of both spectral and colour information. The window size with the equipment used in this study was much smaller than in previous published studies and allows several windows to be placed on a single leaf of a broadleaved plant. Even narrow grassleaved plants could be detected.

The window centres fall on the spectral data line centre location in the images. This line is divided into a number of windows depending on the window size used. The window centres on different spectral data lines are always at the same horizontal locations. This creates a uniform grid of windows in the horizontal direction. Uniformity in the vertical direction is dependent on the movement of the cameras between image pair

acquisition. An example illustration in Figure 26 shows how the square windows are placed on spectral data lines. Distance between the spectral data lines was not perfectly uniform due to inaccuracies in the mechanical gearing system of the camera head and a mismatch between pixel height and camera movement steps. Maximum fluctuations of approximately 2 pixels in colour images were recorded between spectral lines with the test rig setup explained later in chapter 4.0. Red window frames indicate the locations where colour descriptors and spectral data are from plant material only.
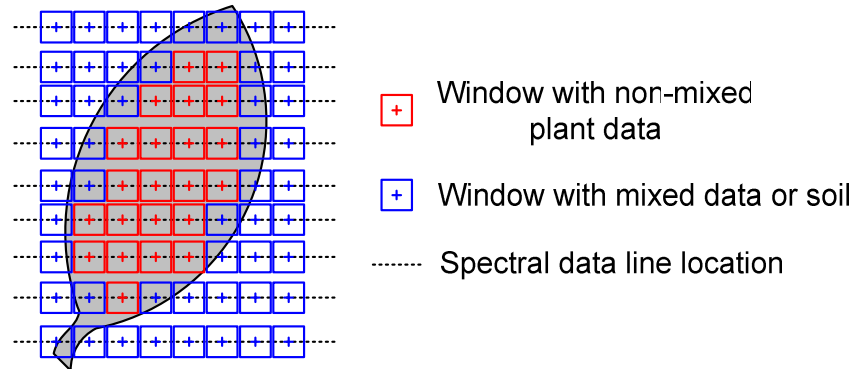
Figure 26 Illustration of window placement principles on spectral data lines

3.3.3 DETECTION OF SEPARATE SMALL PLANT PARTS

The segmentation process rejects objects smaller than 100 pixels in area. This corresponds approximately to an area of 43 mm^2. Also, very thin plant parts may be filtered out in the process of closing and opening the image. By adjusting the smallest object size in the segmentation process some smaller plant parts can be detected but at the same time the noise in the segmentation process has a greater effect.

The limit for the smallest object to be detected in the segmentation is set partially according to the spectrograph spatial resolution. As explained earlier in this chapter, only non-mixed spectral data is recorded. The theoretical spatial resolution of the spectrograph in the used setup was determined to be approximately 0.2 x 1.8 mm. The image pair acquisition interval was based on this spectrograph spatial resolution limit and no image pairs are closer than 2 mm apart. These restrictions basically mean that in order to record non-mixed spectral data the object should be more than twice the size of the spatial resolution limit. This limit is only theoretical and in practice non-mixed spectral data can be recorded rarely for objects smaller than the set area of 100 pixels. An exact measure of an object where one window can be guaranteed to fall on can be calculated from the following equation. All units are in pixels:

$$\begin{bmatrix} width \\ height \end{bmatrix} = \begin{bmatrix} 2(window\_width) - 1 \\ 2(image\_pair\_separation) - 1 \end{bmatrix} \qquad (3.3)$$

Due to different spatial resolutions of the imaging devices, the colour camera can potentially detect smaller plants than the set limit of 100 pixels in area. However, calculation of descriptors may not be meaningful for such small objects.

3.4 PLANT CLASSIFICATION

Previous sections in this chapter have described methods for plant material detection and the calculation and measurement of values to describe the properties of these found objects. This data can then be used to tell the difference between different plant species, i.e. to classify plants.

Previous knowledge of objects beings classified is necessary with almost all methods of classification. This knowledge can be stored in a separate library of data. Some manual intervention is generally needed to create a data library or sample set. This sample set is then used to map certain combinations of feature values to a particular plant and thus creating plant classes. In principle any measurement can then be classified by finding the best match of each measurement to a known plant class.

Multiple discriminant analysis was chosen as the classification method in this project. No single classification method has been shown to be superior, and discriminant analysis has been the choice in many research projects previously. Benefits of the method include low computational load once the classification functions are created and an inherent selection of the best descriptors using the stepwise classification model building method. In addition, plant classification with one non-metric dependent variable (plant class) and multiple independent variables (descriptors) is ideal for the use of multiple discriminant analysis.

The actual implementation of the classification is done in Matlab using functions from a discriminant analysis toolbox created by Vandev (2003). The toolbox offers stepwise linear and quadratic discriminant classification functions and classification performance evaluation tools. Only linear methods were used in this study. Calibration or verification of the mathematical methods in the toolbox is shown in Appendix A.

Use of classification begins by inputting a training set and creating the discriminant functions. The training set includes all the calculated descriptors and a vector describing the class the measurement belongs to. In the training set each plant type was given its own plant class. A stepwise selection process was used to find the most relevant descriptors.

The descriptor with the largest significance was selected as the first variable.

The toolbox functions were altered such that the process of selection was automatic and continued until no variables were added or removed from the model. Originally the selection method in the toolbox was manual and the significance of each variable and the whole model were shown after each step. A significance level of 0.05 was used for inclusion and removal of a variable.

Another addition to the toolbox functions was the checking of multicollinearity of the variables. It is common that plant class spectra are highly similar, especially when narrow band spectral data points are used. This usually means that the data is also collinear. Multicollinearity was checked by calculating a correlation value of a variable against other variables. A variable was accepted to the model only if the correlation value was under 0.9. Highly correlated data can create unwanted bias and degradation of classification results in the model (Hair, Black et al. 2005).

After training or discriminant model building the classification was done using only the set of descriptors selected for the model. A measurement was evaluated with the discriminant functions and the probabilities of a match with the known classes were given. The class with the highest probability was chosen as the class for each measurement.

3.4.1 POSTPROCESSING WITH THE WINDOW METHOD

Classification in the full leaf method was straight forward and no post processing of the classification results was needed. With the window method the classification results were processed using the knowledge of classification results in the neighbourhood of the classified window. This is done because there is often quite substantial variation in the properties of the windows on a single leaf. The leaf area gives a non-uniform reflectance response when investigated at a higher spatial resolution. The filtering method presented here has been iterated through a process of trial and error and thorough visual investigation during the initial testing and has proven to create robust and accurate results.

The post processing of windowed classification results is accomplished by spatial filtering in two layers. The whole filtering scheme is based on a few assumptions:

1. Window classification results have some noise and spatial low pass filtering is a solution
2. A classified window is likely to be surrounded by windows of the same class

The first layer of the filter checks the probabilities of the classes for windows around a 3x3 window neighbourhood of each window and acts as a low pass spatial filter. Probabilities within each class from this 3x3 window area are summed together counting

the centre window probabilities twice. The class with the largest probability sum is selected as a candidate class for the particular window.

The second layer checks the frequency of occurrence of window classes for each class also within a 3x3 neighbourhood of the window and acts as an additional low pass spatial filter using the classification results from the first filtering step. The class with most occurrences in the 3x3 neighbourhood is selected as the final class. In case an equal number of windows are detected for two or more classes, the original window from the first filtering step is the final classification results. This final filtering effectively reduces the number of lonely incorrectly classified windows when surrounded by correctly classified ones.

3.4.2 LEAF BORDER ESTIMATION

Due to the millimetre level spatial resolution of the presented detection system, good estimates of the plant borders in overlapping cases can be given. Classified windows are considered as low-resolution pixels in a higher resolution image so that the leaf border estimation gives a sub-pixel accuracy image of the plants.

The estimation process is illustrated in Figure 27. A binary image of leaves in (a) and classified windows (b) with window locations in the binary image are first needed. For each class, a map is created (c and d), where 2-dimensional bell-shaped normal curves are placed on the window centres and all such shapes are added together within class. The lighter the colour in these class maps, the closer a window centre is. Each pixel in the resulting estimated border image (e) is then given the class number that has that the highest sum in any single class image. The horizontal lines in the class images are due to non-uniform window spacing in horizontal and vertical directions.

(a) Binary Image

(b) Classified Windows

(c) Class 1

(d) Class 2

(e) Estimated borders
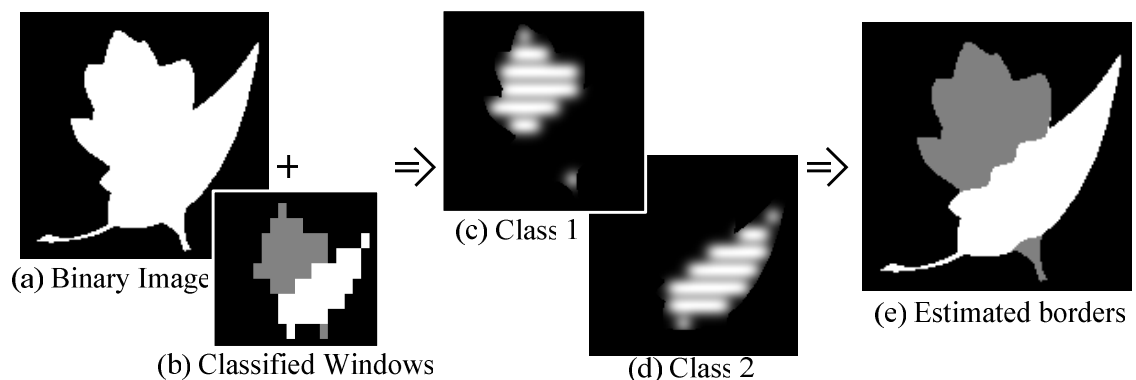
Figure 27 Method of estimating leaf borders from a simulated binary image using classified windows

3.5 DISCUSSION

A novel method for the detection and classification of plants has been presented in this chapter. The method combines the benefits of two imaging technologies, fusing them into a high accuracy sensor for weed detection and classification.

Two data processing schemes for use with the sensor have also been introduced; the full leaf method and the window method. The full leaf method is meant for separate leaf classification using colour, spectral and shape descriptors, while the novel window method is based on classifying small windows in a grid formation. The window method uses colour and spectral data only.

The window method can inherently deal with overlapping plant parts and estimate the location of plants with sub-centimetre accuracy. The method can work well with both broadleaved and grassleaved plants and is not limited to specific plant types. These attributes can be seen as a significant improvement compared to the previously built systems.

Methods of dealing with the increased data flow from two cameras have also been briefly explained. The scope of this project is not to create a system operating in real-time, but to verify the suitability of a new sensor for the application of weed detection and classification. However, no major limitations for a real-time implementation can be seen should correct hardware choices be made.

The presented high resolution estimation of leaf borders is not necessary for the application of chemical or mechanical weed control. However, this method is presented as an innovative way to add value to the raw data from the novel sensor combination. The high resolution capabilities of the system allow new analysis methods to be applied. Leaf border estimation can be advantageous for cases of shape modelling. The output of the estimation could be used with classifiers that have the capability of utilising shape segment information. Such classifiers have been introduced by Petrakis, Diplaros et al. (2002) and Gandhi (2002) and are explained more thoroughly in §2.2.3.

4.0 TEST RIG DESIGN

A test rig was designed for plant identification under laboratory conditions. The basic requirements for the rig were to have a system capable of moving an imaging platform above the plants. This platform will hold colour and spectral imaging devices and a lighting system. The light system is required to block all the external light and provide stable and spectrally known uniform illumination. A description of the test rig design process and an explanation of the choices made is given in this chapter.

4.1 DESIGN EXAMPLES

A number of imaging platforms for weed detection and control purposes have been developed. These research platforms are always customized for a particular purpose and method. Five examples of such platforms can be seen in Vrindts (2000), Hemming and Rath (2001), Åstrand and Baerveldt (2002), Cho, Lee et al. (2002) and Bak and Jakobsen (2004).

When only the imaging systems are considered, the main differences in these platforms are in the selection and setup of imaging and lighting equipment. Common for all the listed platforms is that the imaging device optical axes are perpendicular to the ground, i.e. they are pointing straight down. This is true regardless of the imaging method. In these projects the common method of imaging is conventional RGB colour apart from spectral imaging used by Vrindts (2000). Cameras with different resolutions, lens focal lengths and distances from ground are used, resulting in varying fields of views and spatial imaging accuracies.

Out of the five presented projects only Hemming and Rath (2001) and Åstrand and Baerveldt (2002) use controlled lighting while blocking all the outside light from the imaging area. Halogen lighting is used in both cases. In other projects it is assumed that relatively constant and bright natural light is available.

Cameras are set relatively close to the ground in the platforms while maintaining a width of view reasonable for each particular purpose. The minimum width is in many cases set by the width of the crop row of interest.

4.2 CAMERA STUDY

Camera selection can be of great importance in an application where small reflectance changes are being measured. It is also important to make sure the selected cameras are compatible with all the other hardware. Issues like sensor size, resolution and interfacing

need to be examined.

The imaging platform of the proposed system consists of a colour and a spectral camera. These cameras need to be working together simultaneously with relatively high frame rates. The colour camera specifications are defined mainly by the spatial resolution required. A pixel should reach sub-millimetre resolution. The monochrome camera specifications are defined mainly by the selected spectrograph Imspector V10. This narrows the search of a suitable camera to models offering 2/3 in. (6.6 x 8.8 mm) sensor sizes. Since the cameras need to work seamlessly together, it might be beneficial to have them be from the same manufacturer and with the same interface.

An estimated required field of view width for the cameras is 300 mm. The cameras are designed to be facing straight down at a height of 300-500 mm. The required width of view can be calculated from the following equation.

$$S = \frac{AB}{f} \tag{4.1}$$

In the equation S is the object size, A the distance to object, B the sensor size (height or width) and f the focal length of the lens.

A 2/3 in. sensor is 6.6 mm high and 8.8 mm wide. Solving equation 5.1 for focal length with S=300 mm, A=400 mm and B=8.8 mm gives f=11.7 mm. This is closest to the common lens focal length of 12 mm. This focal length should be used with the spectrograph to achieve the desired field of view.

Affordable colour cameras typically come with a 1/3 in. sensor size with dimensions of 3.6 x 4.8 mm. Estimating the required focal length with the same setup results in f = 6.4 mm. A high resolution lens with a focal length of 8 mm was already available and this requires the height to be 500 mm for a 300 mm width of view. This adjustment still fits in the intended height range so an 8 mm lens will be used with the colour camera.

The next step was to define the required resolution for the colour camera. Given the 1/3" sensor size and an 8 mm lens the field of view at 500 mm is 300 x 225 mm. If only sub-millimetre level is required the sensor resolution should be better than 300 x 225 pixels. Most cameras come with at least a resolution of 640 x 480 pixels, so fulfilling this requirement was not a problem.

In recent years new interface options have become available for machine vision cameras. In many cases a dedicated frame grabber is no longer needed and cameras can be simply plugged in via USB or Firewire ports. Table 2 lists some common camera interface

options with respective data transfer rates.

Table 2 Comparison of camera interface options

Interface	Data transfer rate
Camera Link [*]	2380 Mbit/s
Gigabit Ethernet [*]	1000 Mbit/s
Firewire IEEE 1394b [*]	800 Mbit/s
USB2.0 Hi-Speed	480 Mbit/s
Firewire IEEE 1394	400 Mbit/s
USB1.2	12 Mbit/s

[*] typically requires a dedicated interface card

The three fastest interfaces in the table above typically require a specific interface card to be used. This is especially true with the Camera Link interface. The Camera Link interface also requires special expensive cables to be used. Gigabit Ethernet (GbE) interface offers high data transfer with simple cabling and there are many camera options to choose with this interface. GbE still requires a special interface card but this may change in the near future if this interface standard will be included in the motherboards of common computers. Firewire IEEE 1394b is also a fast interface, but currently there are few cameras utilizing this communication standard.

Most cameras at the moment come with the option of Hi-Speed UBS2.0 or 400 Mbit/s Firewire interfaces. Many Firewire cameras support the DCAM (1394-based Digital Camera Specification) camera interface for easier programming and access to camera settings. Wider selection of camera models and DCAM compatibility were the key reasons to choose IEEE 1394 interface over USB 2.0.

Prosilica (Prosilica 2006) was found to provide Firewire CCD cameras with the required specifications. Models EC650C (colour) and EC1380 (monochrome) were chosen for the test rig with sensor resolutions of 659x493 and 1360x1024 pixels, respectively. The EC1830 has a Sony ICX285 sensor with quantum efficiency shown in Figure 28. The colour camera had a typical spectral sensitivity of an RGB sensor. Since spectral measurements were not taken with the colour camera, its spectral sensitivity was not relevant.
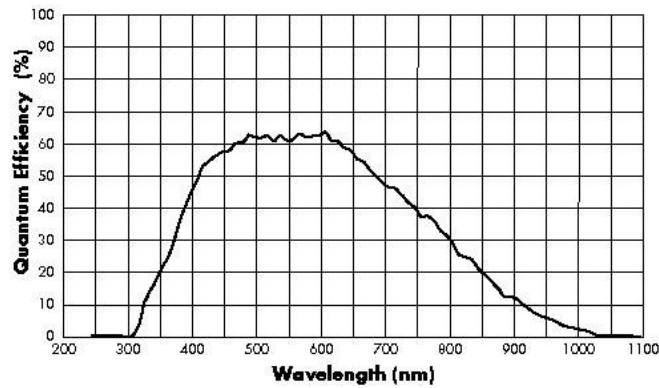
Figure 28 Sony ICX285 CCD sensor Quantum Efficiency (VayTek 2007)

4.3 LIGHT SOURCE STUDY

Selection and design of lighting is always an important issue with imaging devices. Lighting conditions are especially important in spectral reflectance imaging since the cameras measure only the light reflected from the objects. This means that the light source should have relatively uniform spectral irradiance characteristics throughout the wanted spectral range. In this case the spectral range is 400 – 1000 nm.

Heat production and cost of light sources are also issues that need to be considered. Obviously the light sources that produce more light in the infra-red range also produce more heat that may need to be conducted away from the source.

Figure 29 shows spectral irradiance curves of different light sources. Halogen lights give relatively uniform response throughout the VIS – NIR range. Other light sources have distinct peaks in their spectra and provide far less light at the red end of the visible spectrum and NIR ranges. These peaks may become problematic with high spectral resolution spectroscopy applications.
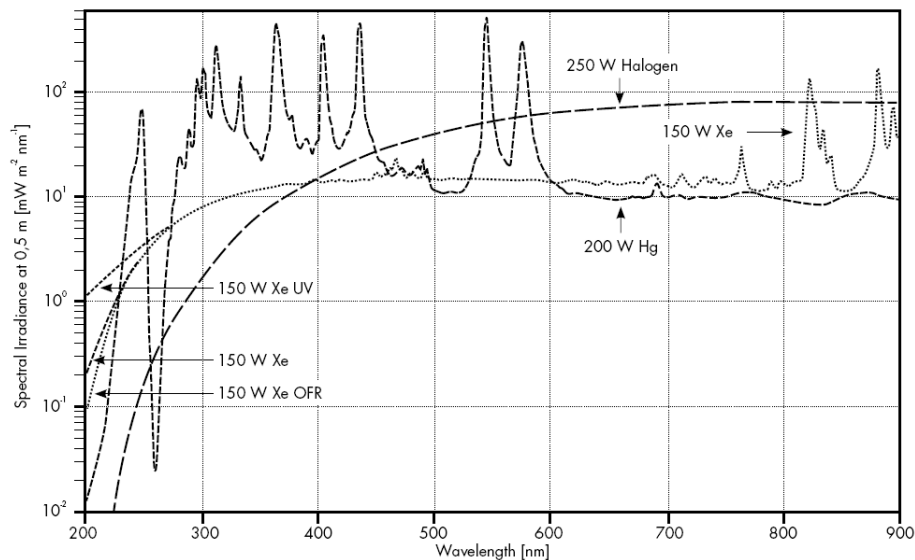
Figure 29 Spectral irradiance of light sources in UV/VIS/NIR range (LOT 2006)

66

Other light sources worth mentioning are white LED (Light Emitting Diode) lights and fluorescent lights. The spectral irradiance curve of a white LED is concentrated on the visible spectrum and they typically have no irradiance in NIR. All common fluorescent light spectra have some distinct peaks in visible spectrum and hardly any irradiance in NIR. The main three fluorescent peaks are produced by the fluorescent process in three different chemical materials in the tubes.

4.3.1 LIGHT SOURCE MEASUREMENTS

Comparative measurements of some light source options were conducted to determine their suitability for imaging in spectral range of 400 – 1000 nm. Spectral response and light intensity were measured and compared.

Figure 30 shows the intensities of some light sources as a function of wavelength. The measurements have been normalized with the CCD sensor spectral efficiency from Figure 28. A perfectly white light source with a uniform intensity throughout the spectral range would be ideal for spectral measurements. Fluorescent light and white LED lights produce near white light observed by human eyes. However, the fluorescent light source shows relatively non-uniform spectra with several peaks in the visible range and no significant intensities at NIR. White LED radiates mostly in the lower part of the visible spectrum but not in the NIR. In fact, human perception of whiteness can be loosely compared to RGB sensor perception of whiteness.
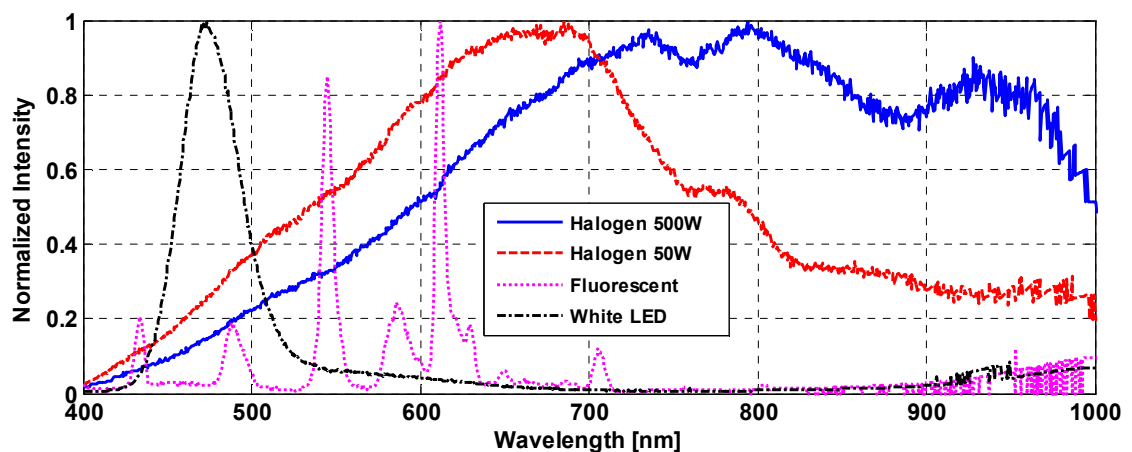
Figure 30 Spectral intensities of different light sources. Values normalized according to the CCD sensor spectral efficiency function.

Halogen lights give the most uniform irradiance in this spectral range. Their irradiance function is strong in the NIR range but not so significant at the lower end of the VIS range. This result is comparable with the spectral irradiance curves presented in Figure 29.

Apart from spectral characteristics, the illuminance of a light source is of importance.

When frame rates are high or objects are moving in the field of view the exposure time needs to be set low. This requires brighter lighting.

Illuminance values of different light sources were also tested with a light meter. The light meter was optimized for the spectral range of a human eye. This means that the light meter measures illuminance values for visible light in the range of approximately 400 – 700 nm only. Measurements were taken directly in front of the light source at a distance of 30 cm.

Accuracy of the light meter measurement tool was 1 Lux but other factors such as positioning of the sensor element created fluctuations of approximately 10% in the final values. The measured light sources were a white LED, a regular long fluorescent light bulb (Polylux XL F58W/835), 50 W halogen spot light and a 500 W halogen flood light. The 50 W halogen directs the light into to a 36° cone and the 500 W flood light was a in metal casing with a back reflector. The measured illuminance values for different light sources are shown in Figure 31. The value in parentheses after the measured value represents how many of the particular lights would be needed to achieve the illuminance power of one 500 W halogen light.
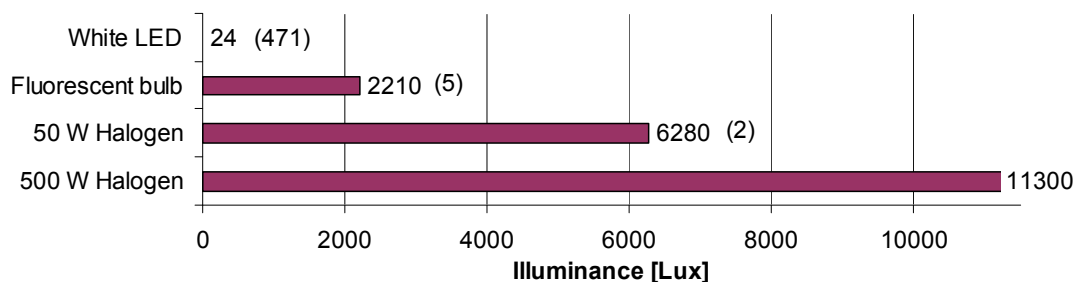
Figure 31 Illuminance [Lux] measurements of light sources at 30 cm distance. The number in parentheses shows the number of lights needed to achieve 500 W halogen illuminance.

4.3.2 CONCLUSIONS

Fluorescent and white LED lights will work well with normal colour imaging, but may not irradiate in the desired wavelengths in narrow band spectral imaging applications. Halogen light sources benefit from low cost and good availability compared to other light sources such as Xenon (Xe) or Mercury (Hg) lamps. Therefore Halogen light sources could be considered the best option for an application of weed detection when spectral imaging is used in the spectral range of 400 – 1000 nm. Halogen lights are also relatively powerful and only a few of them are needed for sufficient illuminance levels.

A combination of white LEDs and Halogen lights could be used to achieve more uniform intensities throughout the whole range of 400 – 1000 nm. The problem of using

LEDs is their relatively poor illuminance. Up to several hundreds would be needed to match the illuminance power of halogen lights. 50W halogen spot lights are chosen as the only light source for the test rig. They are cheap, small and the best compromise between spectral irradiance and heat production compared to 500W halogen flood light.

4.4 ACTUAL RIG DESIGN

All the initial mechanical design of the test rig was done in SolidEdge. Accurate 3D models of all components were made to ensure mechanical compatibility and ease of assembly.

4.4.1 LINEAR MOVEMENT

A linear guide with an 80 cm long movement was modelled in 3D. This guide was to enable the cameras to move over plants. The guide has a flat plate moving along with the motor. The motor in the guide assembly is a unipolar stepper motor attached with rack and pinion drive. An image of the linear guide design is shown in Figure 32.
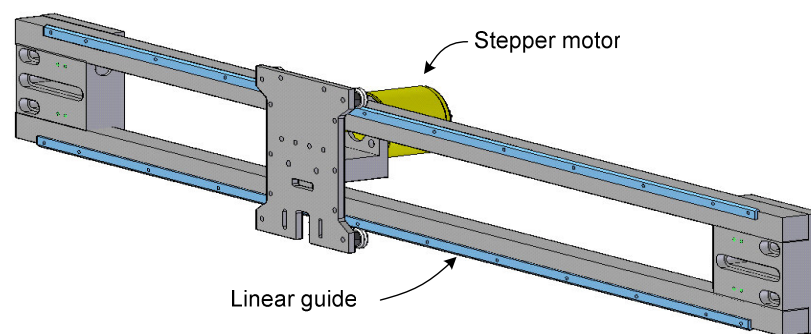
Figure 32   3D model of a linear guide with a stepper motor

To control the motor movement a stepper motor controller and driver were required. In the test rig a microcontroller sends signals to the driver. The stepper motor was driven with full steps of 1.8°, thus the microcontroller was programmed to output four different 4-bit signals in correct order to control the motor movement together with a motor drive board. In the initial design of the test rig the motor control was working individually and did not receive signals from anywhere else.

4.4.2 CAMERAS AND LIGHTING

The cameras and the lighting system need to move together. This means that they both need to be fixed to the linear guide. Additionally the lighting system needs a method to block the outside light and let only the controlled halogen light into the imaging area.

A simple way to block the light is to put a flexible curtain around the imaging area.

Flexibility is needed since the curtain needs to reach the ground but should not disturb it. A round acrylic plate was designed and laser cut in shape to hold the halogen spot lights and give a frame to the curtain to be attached to. The outer edge of the 4 mm thick black acrylic plate was heated and bent down to make the plate more stiff. Halogens lights were fixed to a simple pan and tilt unit for ease of adjustment. This light and curtain combination was then attached to the linear guide. The designed 3D model without the curtain is shown in Figure 33. Two small fans were also installed to the plate to improve ventilation in the area heated with halogen lights.

(a)   (b)
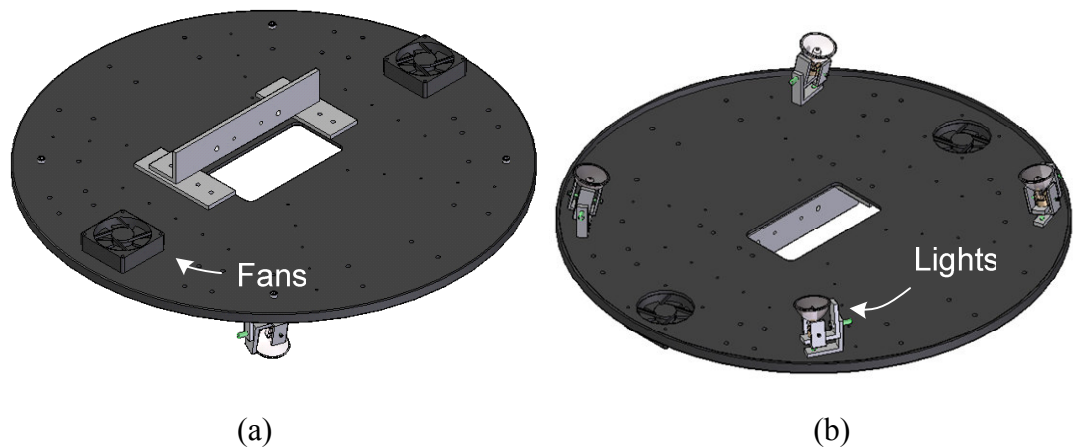
Figure 33 Test rig light assembly (a) from above and (b) from below without a curtain blocking the outside light

An opening was cut to the middle of the plate for the cameras. The cameras were attached next to each other pointing straight down. Due to the slight differences in the camera fields of views the cameras also need to be at different heights. In theory the colour camera should be at a height of 500 mm if the spectrograph is at 400 mm. Due to the spectrograph length a compromise was made to place the spectrograph at a height of approximately 440 mm and the colour camera to 470 mm. This arrangement gives the spectrograph approximately 40 mm wider field of view compared to the colour camera. A 3D model of cameras fixed in place is shown in Figure 34. This assembly was then fixed to the linear guide.
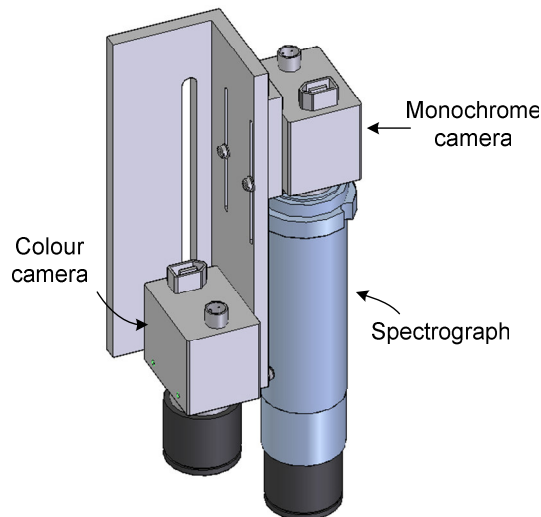
Figure 34 Test rig camera assembly

Since both cameras are required to acquire images simultaneously a hardware trigger signal was generated in the microcontroller. This signal was embedded in the same program with the stepper motor control. The cameras were triggered after a predetermined number of steps. Movement of the rig was stopped while images were acquired and transferred to the computer.

4.4.3 OVERALL STRUCTURE

The actual test rig structure was designed around the linear guide with the cameras and light system attached to it. Aluminium extrusions were selected as the frame structure material. Aluminium extrusions are commonly used in designs for laboratory, the process industry and automation equipment. The frame for the rig was designed in 3D before ordering the extrusions that were cut to length. One of the design criteria was to allow height adjustment of the cameras and lights. The designed frame structure is shown in Figure 35.
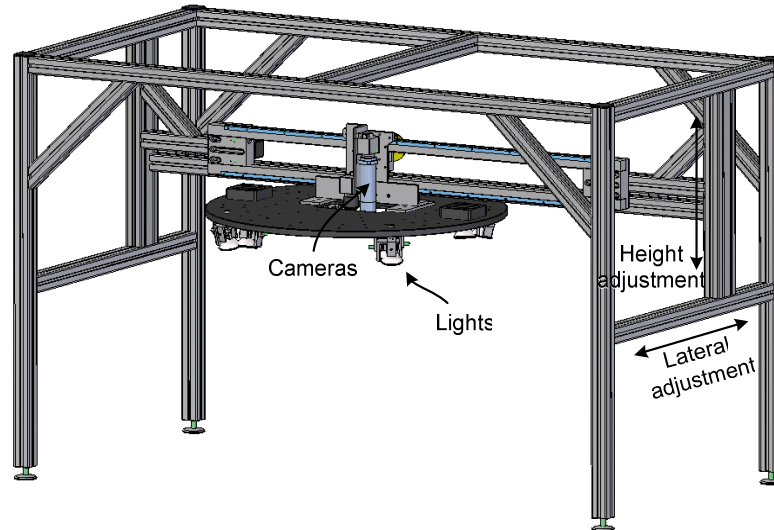
Figure 35 Test rig frame with linear guide, cameras and lights

Since the cameras and lights were on a movable platform the power and data cabling were put in a flexible chain link trunking. The microcontroller and the stepper motor driver were placed with the moving platform and only required power through the trunking. A block diagram of the test rig components and power and data lines is shown in Figure 36.
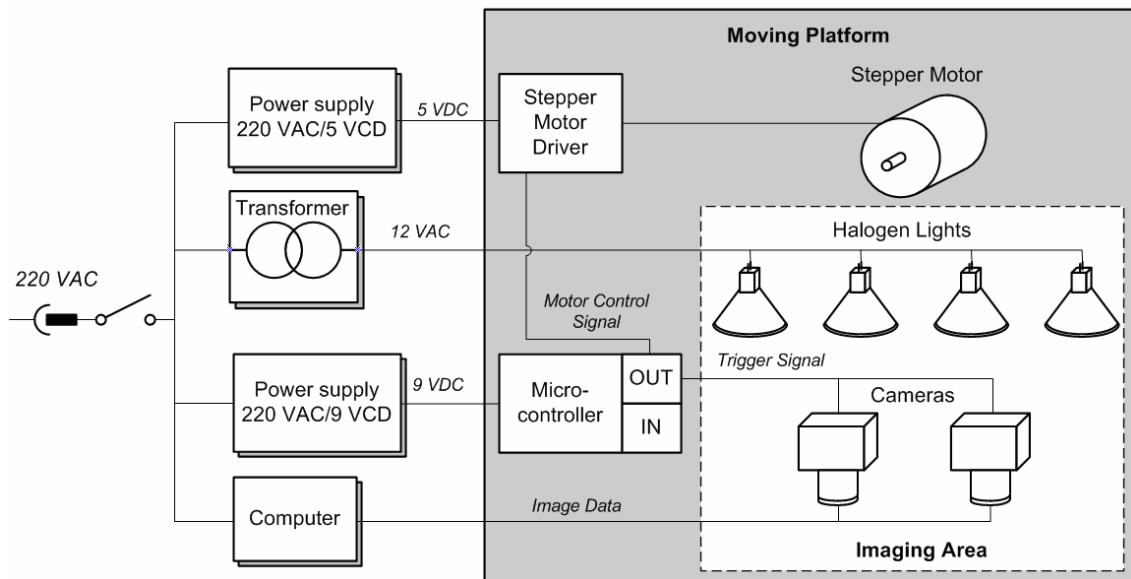
Figure 36 Block diagram of the test rig systems

Another conceptual image of the test rig and the camera system is shown in Figure 37. The sketch shows how the colour and spectral camera fields of view are overlapped and how the coordinate system is placed. Approximate dimensions of the colour camera field of view are also shown. The sketch also shows how the spectral camera disperses a line in the colour image into spectral components. Vertical image axis of the spectral image

shows the spectral content, while the horizontal image axis maintains the spatial information. The linear guide moves the camera system along the y-axis only.
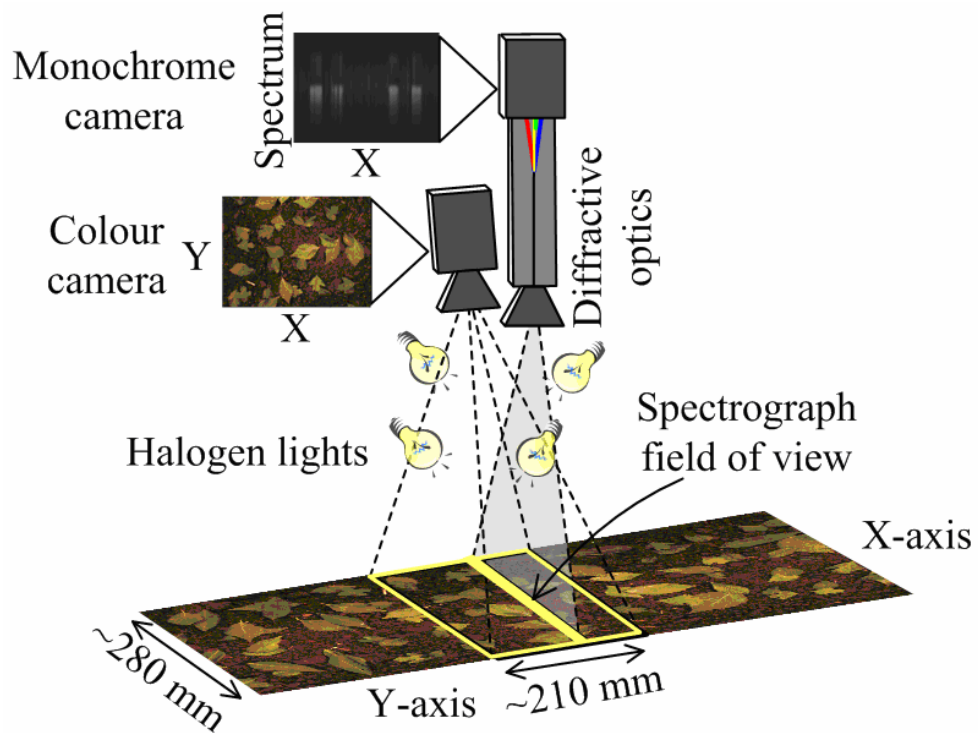
Figure 37 Sketch of the test rig camera setup and coordinate system

4.4.4 DESIGN ASPECTS FOR A MOBILE ROBOTICS APPROACH

The ultimate goal with the plant classification system was to work autonomously on a mobile platform. The current test rig is limited to data collection in a small area in laboratory conditions, and great changes would need to be implemented for the system to operate autonomously in field conditions.

The current test rig allows only linear movement of the imaging devices. A major change with the mobile platform would be to attach all of the equipment to the main structure or frame. This frame would then navigate in the field with a help of additional sensors such as cameras, GPS, inertial units or odometric devices.

Extensive research on various mobile platforms for weeding purposes has been done (Åstrand and Baerveldt 2002; Søgaard and Lund 2006). There are still several obstacles to overcome before commercial emergence of these platforms. This project concentrates on the recognition and classification of plants. The test rig system is therefore designed for laboratory and limited outdoor use only.

5.0 LEAF CLASSIFICATION IN LABORATORY CONDITIONS

Laboratory tests were conducted to check the performance of the developed system and algorithms in a controlled environment. Laboratory conditions can be adjusted to exclude external influences such as unknown plants and their positions, foreign objects in field of view and wind. In a real outdoor field environment this may not be possible. The following sections explain the calibration of the equipment, actual test procedures and results from the laboratory testing.

5.1 CALIBRATION

Calibration of the colour camera, the spectrometer, the halogen lights and the linear movement was needed in order to obtain useful data from the system. The calibration of the cameras and lights is unique for each setup and should be done every time any hardware adjustments are made. The following sections explain the procedures of calibration for each component.

5.1.1 LINEAR MOVEMENT CALIBRATION

The linear movement system consists of a unipolar stepper motor on a rack and pinion drive. The stepper motor driver is a combination of a microcontroller board and a custom built driver commanding the stepper motor with full step increments. A full step with the motor is 1.8°. The rack has four teeth/10 mm and the gear head 12 teeth/360°. From these values, a movement of 0.15 mm/step can be calculated.

5.1.2 SPECTRAL CALIBRATION OF THE SPECTROGRAPH

The spectrograph needs to be spectrally calibrated so that measurements with it can be related to actual wavelengths. The spectrometer Imspector V10 has a nominal spectral range of 400 – 1000 nm. Calibration finds the relationship between the wavelengths and their spectral positions in the image (horizontal axis or image rows). Before calibration could begin the spectrograph was aligned with the imaging sensor axis so that all spectral lines were parallel to the image rows of the camera sensor.

Spectral calibration was done using common fluorescent lights as a light source. All fluorescent light spectra have several easily resolvable peaks that can be used for spectral calibration. A Spectralon reference tile was illuminated by the light source and the reflected light was imaged with the spectrograph.

The imager used with the spectrograph was a Prosilica EC1380. The resolution of the sensor was 1360x1024 pixels, giving 1024 data points along the spectral axis. The

theoretical spectral resolution is then (1000-400 nm)/1024 pixel = 0.59 nm/pixels, which is almost 7 times better than the nominal spectral resolution of four nanometres given by manufacturer's specifications. This does not however mean that the equipment can reliably be used to separate wavelength differences smaller than the nominal spectral resolution.

Figure 38 shows the spectral reflectance from the Spectralon tile with only fluorescent light as an illumination source. Calibration data points are overlaid in the figure at their respective spectral peaks. The relationship between the wavelengths and the data point locations (y-axis values in the spectral images) is estimated by fitting a line through the points. A linear relationship is expected. Such a fitted line is shown in Figure 39, with the function used to calculate the wavelength for each data point.
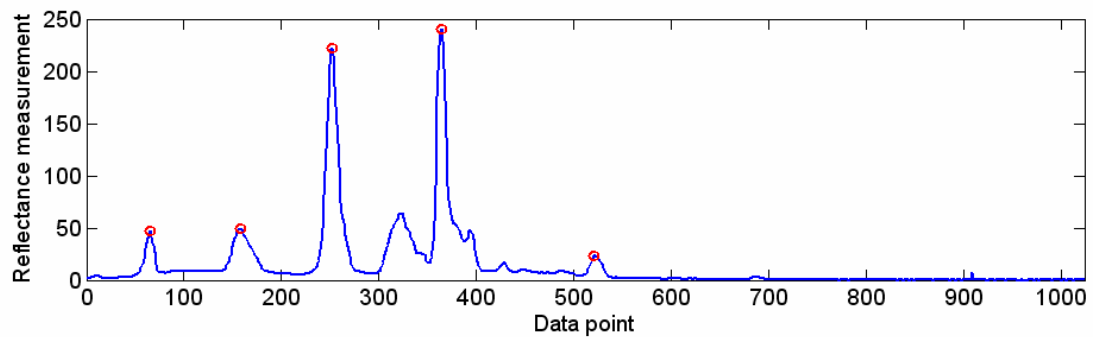
Figure 38 Reflectance measurement of fluorescent light with overlaid calibration data points.
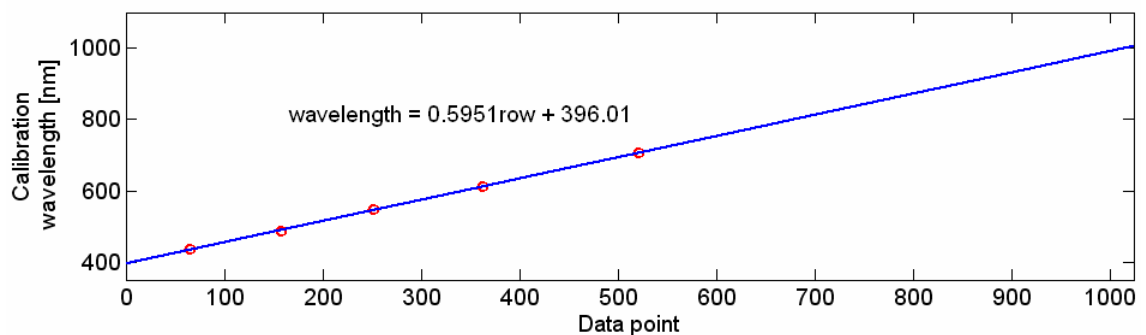
Figure 39 Spectrograph calibration curve with overlaid fluorescent light spectral peaks and the calibration function

According to the estimated function, the spectral range of the spectrograph is 397 – 1006nm. This is very close to the range given in the manufacturer's specifications. Table 3 shows the spectral calibration points and their error with respect to the estimated relationship function. All the calibration points are within a specified spectral resolution of 4 nm.

Table 3 Spectral calibration points and error to the estimated relationship function

Fluorescent peak [nm]	Respective data point	Error to linear estimation [nm]
436.6	59	-1.8
487.7	147	0.6
546.5	243	0.2
611.6	356	3.8
707.0	502	-2.8

The spectral calibration did not take absolute reflectance intensities into consideration. The final application has controlled lighting conditions and only relative differences in reflectance values are of interest. Also, "smile" and "keystone" effects were not corrected during calibration. Smile means curvature distortion of the spectral image, while keystone effect distorts a rectangular focal plane into a trapezoid. A more thorough and accurate spectral calibration process is described by Lawrence, Park et al. (2003), including calibration for spatial dimensions, spectral wavelengths and percent reflectance values.

5.1.3 SPECTRAL DATA SPATIAL RESOLUTION

The Imspector spectrograph, just like any other spectrograph or imaging device, has a limited spatial resolution. A spectral measurement at a specific location with the Imspector gives average spectral properties of that spatial area. When this area is spectrally homogeneous the measurement is considered unmixed. If the area consists of objects with different spectral properties the signal from it is a mix between these properties. Area proportions of different spectral regions govern how much of each property is present in the overall signal if a linear spectral mixing model is assumed. Unmixing such signals often requires complex algorithms (Keshava 2003).

The spatial resolution of Imspectror is determined by four factors – the objective lens focal length, distance to an object, the input slit width and length and the imaging sensor resolution, as labelled in Figure 10. The spectrograph spatial resolution in horizontal x-axis direction is mainly governed by the imaging sensor resolution and the vertical y-axis spatial resolution by the input slit width. Spectral resolution should not be mixed with spatial resolution as it is a separate issue determined by different factors.

The Imspector optics have a magnification of 1, and the theoretical length (L_i) and width (W_i) of the spectral line (physical field of view of the Imspector) can be calculated with the following equations

$$L_i = L_s D / f \tag{5.1}$$

$$W_i = W_s D / f \tag{5.2}$$

where L_s is the input slit length, W_s the input slit width, D the distance to the target and f the lens focal length (Specim 1997). Theoretically the spectrograph therefore has a spectral line 322.7 mm long and 1.8 mm wide in the given setup; L_s = 8.8 mm, W_s = 50 μm, D = 440 mm and f = 12 mm.

To determine the true spatial resolution of the spectrograph, a pattern shown in Figure 40 was used. The spectrograph field of view was scanned across the pattern in the vertical y-axis direction with 0.15 mm increments to find the vertical spatial resolution. A spectral line at approximately 650 nm was recorded from each measurement and these lines were put together to form the test data image shown in Figure 41. In this image the data from separate images are shown in the vertical measurement axis and the horizontal axis has the spatial data. This test image shows a portion of the test pattern in Figure 40 imaged with the spectrograph. The axes in Figure 41 are not to the same scale as the test pattern in Figure 40.
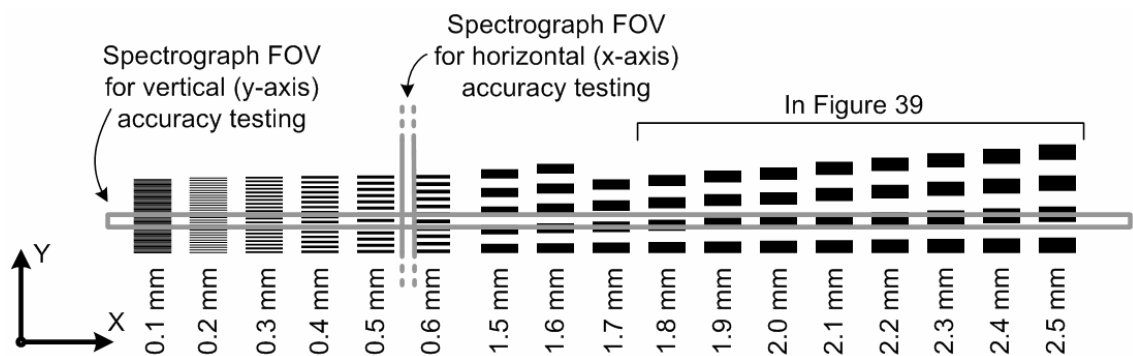
Figure 40 A pattern for spectrograph spatial resolution or sampling accuracy testing. Millimetre values represent the line and gap width. The figure is not in scale.

The vertical spatial resolution or spatial sampling width can be determined from Figure 41 by judging when the line separation is large enough for the spectrograph to measure only white background paper in the gaps of lines. This happens when line separation is 2.5 mm and is indicated in Figure 41 by showing how white background emerges between the pattern lines in the measurements. When both black lines and white paper background are in the field of view, the output is a mixture of these signals proportionate to the area coverage of white and black. This mixing can be seen as shades of grey in Figure 41. The resulting resolution of 2.5 mm is somewhat lower than the theoretical minimum of 1.8 mm.
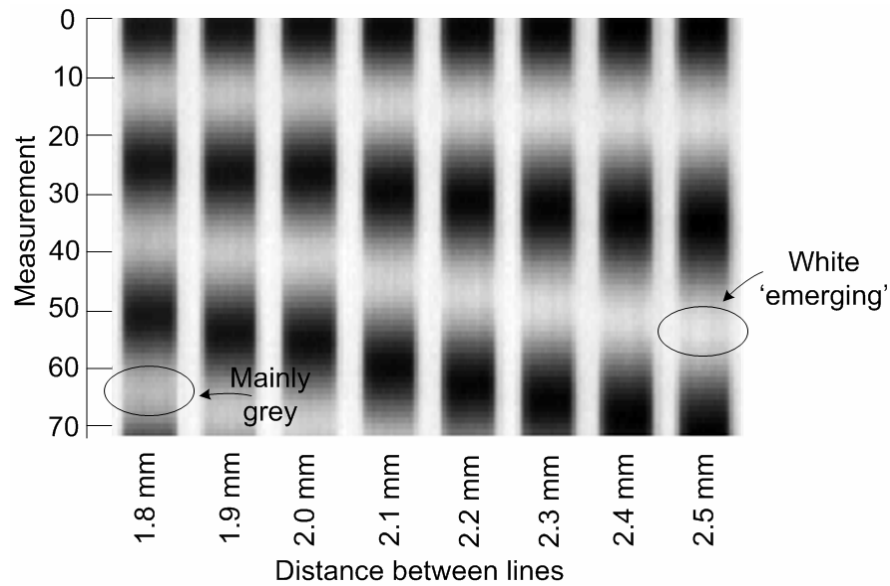
Figure 41 Series of images measured with a spectrograph along part of the spatial resolution test pattern

The horizontal spatial sampling accuracy depends highly on the imaging sensor resolution and fore lens focus. The same test pattern in Figure 40 was used to determine the smallest horizontal line gap size visible in spectral images. This time only one image per line size was needed and the smallest visible gap between lines was 0.3 mm. These results give us dimensions of the smallest possible homogeneous area that the spectrograph can image without spectral mixing with the given setup. The dimensions are 0.3 x 2.5 mm.

During the laboratory testing, the internal focus of the spectrograph was adjusted such that the resolution W_i was 6 mm. The cameras were also moved by 6 mm along the linear guide between each image pair. The lower resolution was selected to reduce the amount of data when relatively large leaves were imaged.

5.1.4 COLOUR AND SPECTRAL DATA REGISTRATION

The detection system uses data from two separate cameras. In order to be able to combine the information from these cameras the data needs to be registered or fused together. Registration is a process where data from different sources is transformed into one coordinate system. As shown in Figure 37, the colour camera and spectral camera fields of view overlap. The precise location of the spectral FOV in the colour image needed to be defined for accurate data fusion.

The first step to find the spectral FOV location in the colour image was to make sure that the cameras were aligned properly. A test pattern in Figure 42 was used to verify that the spectral line was parallel to the colour camera x-axis. The pattern was placed under the colour camera horizontally and the spectral camera was rotated and fixed in a position

where the spectral imaging line was precisely along the central line of the pattern. This happens when there are uniform thin vertical black lines in the spectral image. If the spectrograph is rotated relative to the pattern the black lines in the spectral image do not have a uniform thickness. Also, when the spectral imaging line is not in the centre of the pattern the black lines grow thicker. When the spectrograph is rotated perfectly relative to the colour camera x-axis and the test pattern, the location of the centre line of the pattern can be recorded in the colour image.
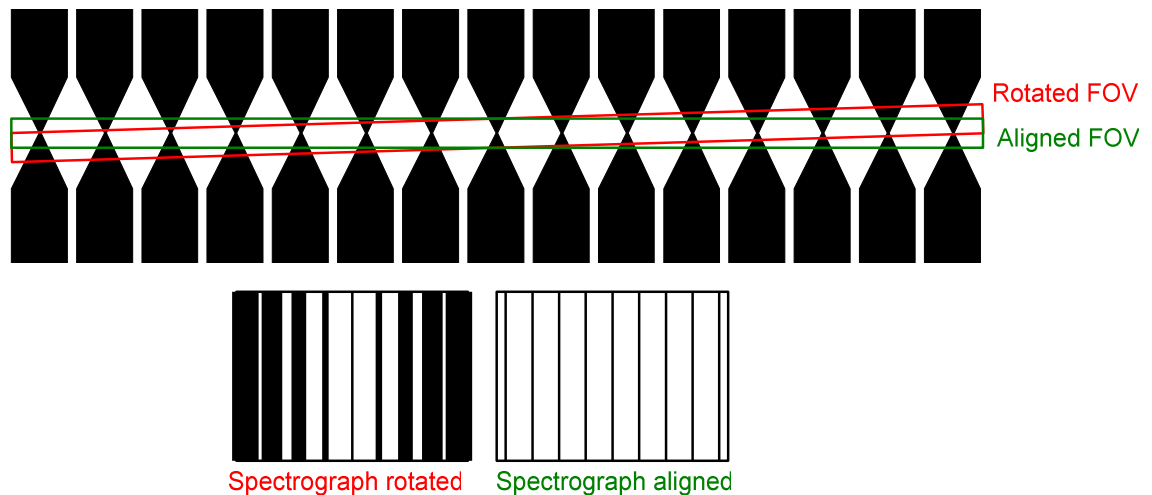
Figure 42 Alignment pattern and examples of rotated and aligned spectrograph FOV output.

The spectrograph imaging line FOV width in the setup is greater than the colour camera FOV. The left and right sides of the spectral image can then be cropped to match the colour camera image horizontal FOV. Since in this case the colour camera sensor resolution is smaller than the spectrograph imaging sensor resolution the cropping should be done in spectral images or some overall spatial resolution of the system will be lost.

The final step in registration after camera alignment, recording the spectral imaging line location in the colour image and cropping the spectral image was to interpolate the cropped spectral image x-axis to match the x-axis locations of the colour image using an imresize-function from Matlab Image Processing Toolbox. This ensures that each data column in the processed spectral image corresponds to the same column number in the colour image.

5.1.5 EFFECT OF A SPECTRALLY OPTIMIZED LENS

Each optical element between the object and the camera sensor can only make the ideal image worse. Therefore, when working at the sensitivity limits of a camera, the best optical equipment should be used when possible. Regular lenses work rather well for applications

in the visual spectrum, but start to dramatically lose their transmittance at the near infrared region. The near infrared region has been considered especially important when plants are being imaged.

Spectrally optimized lenses for the sensitivity band of silicon sensors (400-1000 nm) have close to uniform transmittance curve throughout the whole band compared to a regular lens as shown in Figure 43. These optimized lenses will let significantly more light into the sensor at the near infrared end of the spectrum compared to normal lenses. The measured differences in the lenses are illustrated in Figure 44, and the results match well with the specification data. The improvement gained with an optimized lens proves valuable in spectral imaging especially in the near-infrared region.
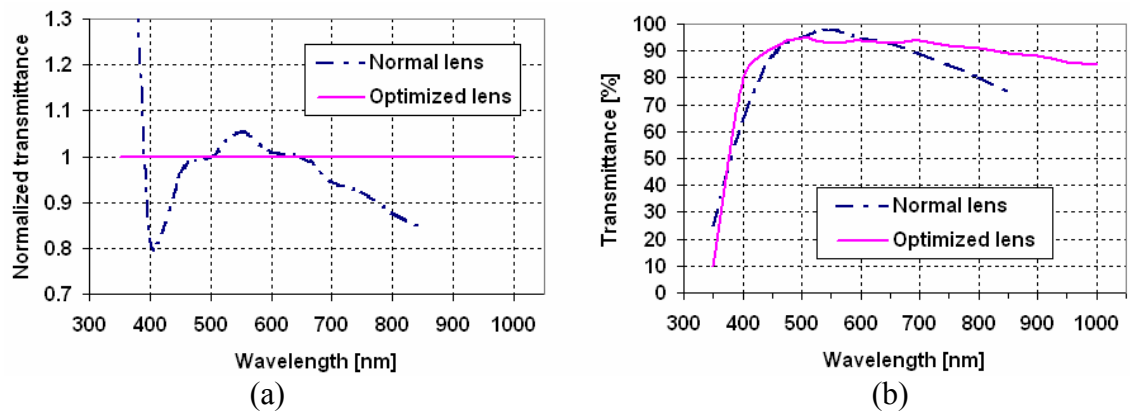
Figure 43 Comparison of normal and 400 - 1000 nm optimized lens transmittance from specification data. a) reproduced from (EdmundOptics 2005)
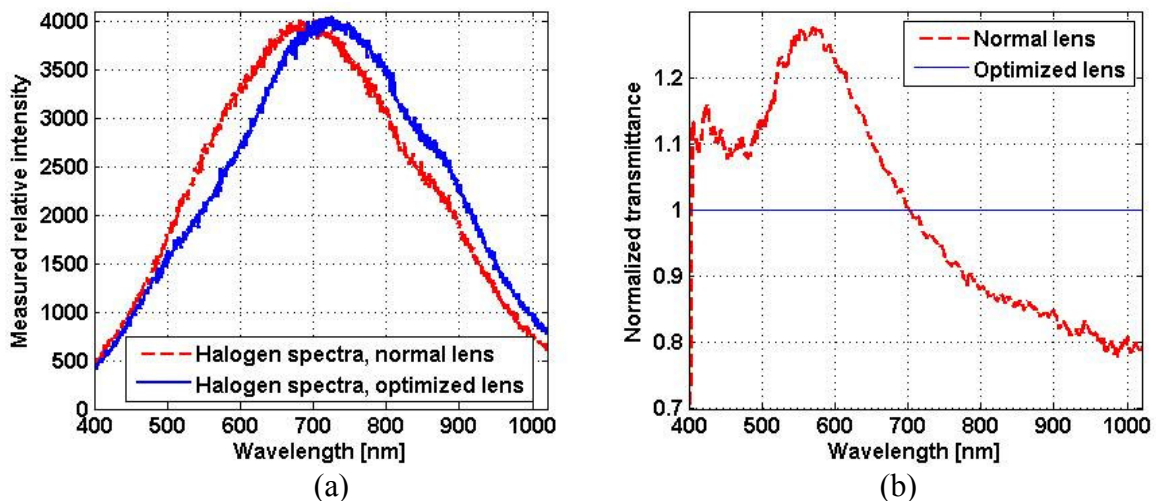
Figure 44 Comparison of measured lens transmittance between normal and 400 - 1000 nm optimized lens. a) Measured intensities as a function of wavelength and b) the lens transmittance figures normalised by the optimised lens values

5.1.6 LIGHTING CALIBRATION

Four halogen spot lights provided the lighting for the imaging area at a height of

approximately 400 mm from the ground. The lights were positioned around a circle of 310 mm radius, 90° apart from each other. Each light was pointing roughly to the centre of the imaging area. The lighting setup did not produce a perfectly uniform light intensity throughout the whole imaging area. This could be accounted for by using a simple calibration technique based on full frame intensity images.

To calibrate the system for these unchanging intensity variations a set of three images of white paper were taken with the colour camera. The white paper filled the entire FOV. The average value for each pixel from the three images was calculated and each RGB channel was separated with their intensity values normalised to have values of 0 to 1. The red channel calibration image is shown in Figure 45. A colour image with uniform lighting was then achieved by dividing the acquired image with the calibration image values. This operation was done for each RGB channel individually.
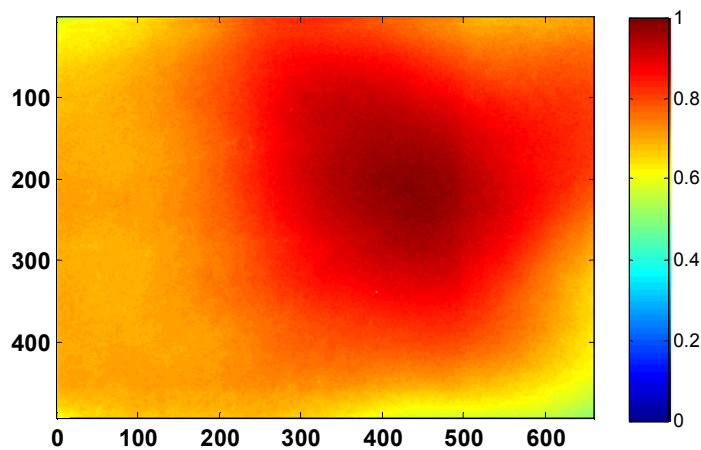
Figure 45 Lighting calibration image for red channel showing relative intensity

A slightly different approach was used when calibrating the lighting intensity in spectral images. The spectrum of the Spectralon reflectance plate was first measured at each horizontal location and a wavelength that produced the highest intensity values (a row in the spectral image) was chosen as calibration values. These calibration values were normalized from 0 to 1. Each acquired spectral image row was then divided by these calibration values to calibrate for uniform lighting conditions.

5.2 DATA COLLECTION

A selection of outdoor plant leaves were gathered and imaged from around the laboratory building. Six plant classes were used in the laboratory tests. These were Hawthorn (*Crataegus monogyna*), Rugosa Rose (*Rosa rugosa*), Ash (*Fraxinus excelsior*), Silver Birch (*Betula pendula*), Berberis (specific type unknown) and Blackthorn (*Prunus spinosa*). These plants were 1-6 respectively. This group of plants was selected for the study simply

because the leaves were readily available in the vicinity of the laboratory and could be imaged immediately after being picked.

The selected plants are either trees or shrubs. No actual crop or weed plants were used at this stage of the study. The system needed to be able to correctly classify any selection of plant leaves and therefore it was not necessary to be limited to actual crop or weed plants.

The collected leaves were placed flat on the bare soil platform in random order and orientation such that no leaves were overlapping. Leaves were not forced flat and the natural curvature of the leaves remained intact. The number of leaves collected for each class is shown in Table 4.

Table 4 Number of leaf samples and distribution of training and sample sets in laboratory testing for each plant class

Class	Leaves in total	Whole leaf method[1]		Window method[2]	
		Training	Sample	Training	Sample
1 Hawthorn	114	30	84	399	963
2 Rugosa Rose	106	30	76	607	1602
3 Ash	72	30	42	1180	1401
4 Silver Birch	60	30	30	1778	1578
5 Berberis	81	30	51	408	474
6 Blackthorn	108	30	78	184	581

[1] Includes colour, shape and spectral descriptors
[2] Includes colour and spectral descriptors

The classification training set was chosen from the first 30 leaves imaged for each class. With the window method the same 30 leaves were used to extract small windows and the number of training set measurements was dependent on the leaf size as more windows can fit on bigger leaves.

The imaging and calculation of descriptors was done as described in chapter 3.0 using a 6 mm camera movement between image pairs. Six imaging runs were needed with the test rig to gather the data listed in Table 4. All of the imaging conditions were identical during the six test runs. An additional test run was taken to further evaluate the window method with overlapping leaves. During this test leaves from classes 1-4 were placed on the soil platform partially on top of each other in random orientations. One blob consisted always of two or three leaves and it was ensured that all possible combinations of leaf pairs were available at least twice. This included overlaps within a single plant class. Combined colour images for each of the seven test runs are shown in Figure 46. Each colour image shown consists of 130 separate images registered together.

Class 1, Hawthorn

Class 2, Rugosa Rose

Class 3, Ash

Class 4, Silver Birch

Class 5, Berberis

Class 6, Blackthorn

Overlapping leaves from classes 1-4
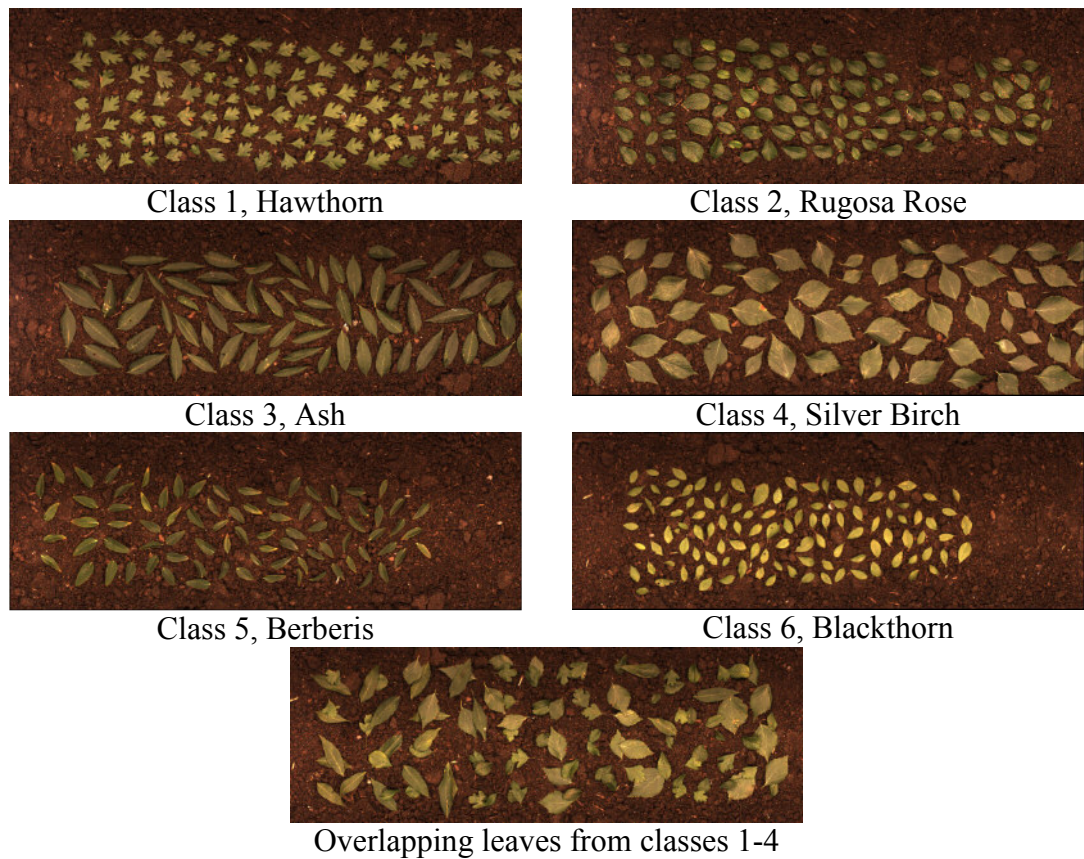
Figure 46 Combined colour images from seven laboratory test runs

The average normalized spectral reflectance curves for each class are shown in Figure 47. This data was taken from the full leaf method. The between-class differences in these reflectance values are really small and can give an idea on how difficult the task of classification with this data can be. Some distinct differences between classes can be seen especially in class six, but if several individual measurements from each class were plotted on this figure they would be hard to separate.
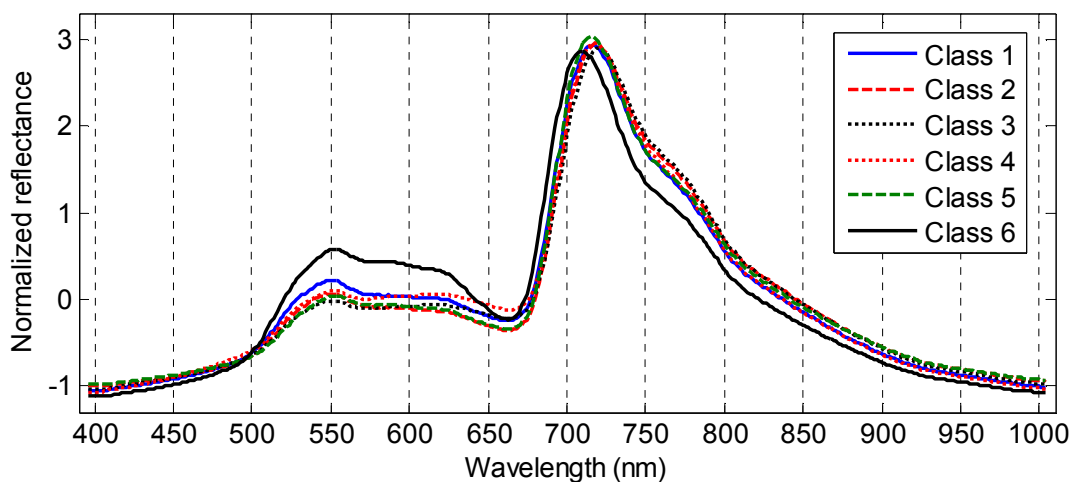
Figure 47 Average normalized spectral reflectance for 6 plant classes

Figure 48 shows an example of how plant classes are separated with certain descriptors. The descriptors used in the figure are taken from the full leaf method and with a spectral resolution of 11.3 nm. Using only four descriptors of eccentricity, 257 nm, 314 nm and 219 nm, the classification was 89.0% accurate. From these two plots it can be seen that class 6 is easily separated from other classes in both plots. Classes 2 and 4 are well separated by descriptors 219 nm and 314 nm but slightly overlapping if eccentricity and 314 nm descriptors are used.
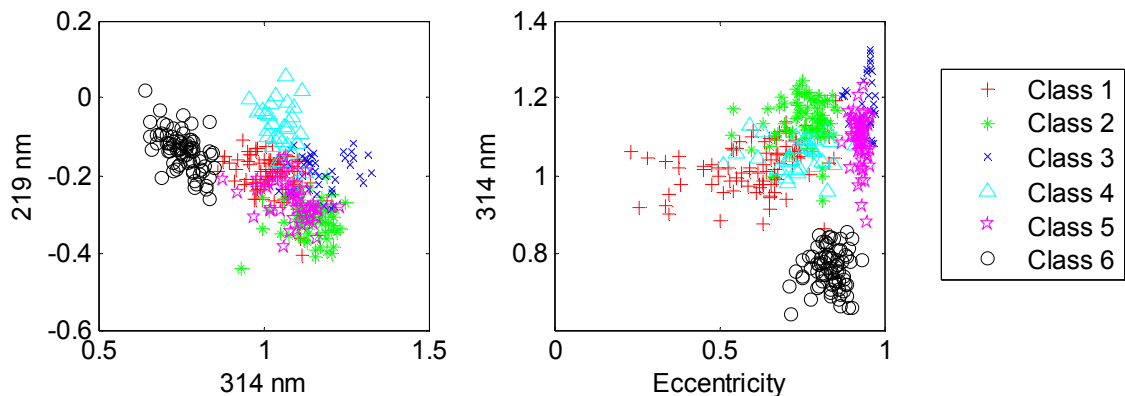
Figure 48 Example of differences between classes with three descriptors. Data taken with 11.3nm spectral resolution from full leaves

5.3 RESULTS

The results from the laboratory tests are shown and analysed in this section. Results from two different tests are first introduced. These tests classified data collected either from full leaves or windows on leaves. In the first analysis with full and separated leaves, use of descriptors from shape, colour and spectrum are compared. In the window classification analysis use of shape was not done and results are shown with separated and overlapping leaves.

Leaf border estimation analysis with window method is done to provide results for comparison with previous studies. Ideal descriptors in both full leaf and window method are also investigated. The analysis method of descriptor occurrence frequency was chosen for quantifying the importance of the descriptors to allow comparison with results from previous studies.

5.3.1 WHOLE LEAF METHOD

A number of different classification approaches were evaluated for the whole leaf method. The set of training descriptors available for the classification algorithm and the spectral data point resolution were varied. The same descriptor sets were used in training and in actual performance evaluation with a sample set. Five different descriptor sets were

constructed from three sets of shape, colour and spectral descriptors. Spectral data point resolution was changed from 13.1 nm to 72.6 nm in nine steps. This was done artificially by averaging together different numbers of data points along the spectral axis for different spectral resolutions.

Various sets of descriptors were used since the stepwise selection process in multivariate discriminant analysis is greatly affected by the available descriptors. Also, it is important to find if the combination of information from colour and spectral cameras is beneficial for classification purposes. The spectral point resolution was changed to find what effect this would have and if lower spectral resolution could be used without compromising the classification accuracy. Lower spectral resolution requirements could potentially allow the use of even lower cost cameras.

The classification performance with five sets of descriptors and varying spectral resolutions with the whole leaf method is shown in Table 5. The percentages shown are average values over all six classes and the best case for each descriptor set is in bold. Spectral resolution is not applicable for the first descriptor set of colour and shape and only one performance value can be calculated.

Table 5 Average classification results with various combinations of descriptors and spectral resolutions with the full leaf method

Spectral resolution	Descriptor set				
	Colour and shape	Spectrum	Colour and spectrum	Shape, colour and spectrum	Shape and spectrum
13.1 nm		90.2%	92.1%	96.4%	96.8%
15.5 nm		91.3%	91.2%	97.2%	97.2%
17.9 nm		91.4%	91.4%	94.3%	94.3%
20.2 nm		89.5%	89.5%	97.1%	97.1%
22.6 nm	**79.7%**	91.9%	**92.5%**	**97.6%**	**98.0%**
36.9 nm		89.5%	87.4%	97.5%	97.5%
48.8 nm		90.2%	89.7%	96.9%	96.2%
60.7 nm		**92.4%**	89.0%	97.2%	97.1%
72.6 nm		90.5%	84.8%	96.8%	96.8%

Figure 49 shows more detailed information for the best cases of each descriptor set. The classification results for each plant class and different descriptor sets are visualized together. It can be seen that performance is best in all classes when shape and spectrum were used. The performance is only slightly lower when all descriptors including colour were used. This minor difference was caused by the stepwise descriptor selection process and shows that using all variables may not always produce the best results.
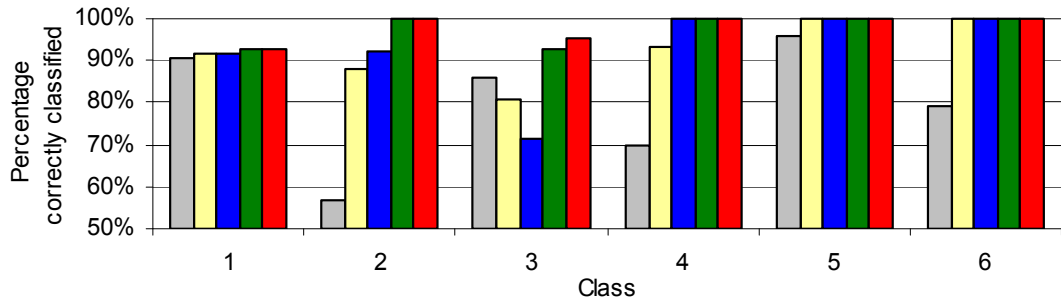
Figure 49 Performances of full leaf classification in each class with ▢ RGB and shape, ▢ spectrum only, ■ RGB and spectrum, ■ RGB, shape and spectrum and ■ shape and spectrum descriptors

With the three best performing descriptor set selections, classes 4-6 are consistently classified perfectly within the given sample measurements. Class 3 seems to be the most difficult to classify correctly and class 1 classification performance is nearly independent of the descriptor set selected.

Table 6 shows more detailed information from the best performing cases for each descriptor set in the form of classification error matrices. The diagonal elements in the matrices show the correctly classified cases. Each column adds up to 100% (with the exception of errors from rounding the decimals). Non-diagonal elements on each column represent the percentage and distribution of misclassified cases for a particular class. The *Total*-column at the end of each row sums up the percentages on each row, and represents how prone the classifier is to misclassifying measurements to a particular class. A value above 100% in this column indicates that samples are often misclassified to the particular class (false positives).

Table 6 Classification error matrices for the best performing cases within each descriptor sets

		Colour and shape							Spectrum only (λ_{res}=60.7nm)						
Class		1	2	3	4	5	6	Total	1	2	3	4	5	6	Total
Classified as (%)	1	**90.5**	2.6	0	0	0	3.9	97	**91.7**	7.9	0	3.3	0	0	102.9
	2	0	**56.6**	0	0	0	0	56.6	6	**88.2**	16.7	0	0	0	110.9
	3	0	0	**85.7**	3.3	2	0	91	1.2	3.9	**81**	3.3	0	0	89.4
	4	0	1.3	0	**70**	0	0	71.3	1.2	0	0	**93.3**	0	0	94.5
	5	3.6	11.8	14.3	16.7	**96.1**	16.9	159.4	0	0	2.4	0	**100**	0	102.4
	6	6	27.6	0	10	2	**79.2**	124.8	0	0	0	0	0	**100**	100

		Colour and spectrum (λ_{res}=22.6 nm)							Colour, shape and spectrum (λ_{res}=22.6 nm)						
Class		1	2	3	4	5	6	Total	1	2	3	4	5	6	Total
Classified as (%)	1	**91.7**	7.9	11.9	0	0	0	111.5	**92.9**	0	0	0	0	0	92.9
	2	6	**92.1**	14.3	0	0	0	112.4	7.1	**100**	4.8	0	0	0	111.9
	3	0	0	**71.4**	0	0	0	71.4	0	0	**92.9**	0	0	0	92.9
	4	2.4	0	2.4	**100**	0	0	104.8	0	0	0	**100**	0	0	100
	5	0	0	0	0	**100**	0	100	0	0	2.4	0	**100**	0	102.4
	6	0	0	0	0	0	**100**	100	0	0	0	0	0	**100**	100

		Shape and spectrum (λ_{res}=22.6 nm)						
Class		1	2	3	4	5	6	Total
Classified as (%)	1	**92.9**	0	0	0	0	0	92.9
	2	7.1	**100**	4.8	0	0	0	111.9
	3	0	0	**95.2**	0	0	0	95.2
	4	0	0	0	**100**	0	0	100
	5	0	0	0	0	**100**	0	100
	6	0	0	0	0	0	**100**	100

5.3.2 WINDOW METHOD

Classification with the window method was done using colour and spectral descriptors. The performance of seven different spectral resolutions was tested, ranging from 6.0 nm to 20.2 nm. This resolution range begins near the spectrograph limits and extends to a range where classification performance was noted to degrade. The different spectral resolutions were calculated by sub-sampling the spectral images in the spectral axis by adding together different numbers of data points along this axis.

The number of cases in classification training and sample sets were shown earlier in Table 4. Since the number of windows per leaf is dependent on the size and shape of the leaf, the number of windows per plant class changes significantly. These differences may bias the classification results, increasing the number of measurements classified to classes with higher number of training cases used. To find out the effect of uneven number of training cases per class the classification was done using windows from 30 leafs (non-

sampled training data) and also using randomly sampled windows (sampled training data) totalling to 184 training windows per class, as this was the smallest number of windows for 30 leaves in any plant class used in this study.

The performance figures of the average classification with varying spectral resolution are shown in Table 7. The best average performance was recorded at 10.7 nm, and classification accuracies for each class with this spectral resolution are shown in Figure 50. With sampled results the number of training cases was made equal for all classes by randomly sampling within the full training set data. The figure also shows the performance of classified windows after spatial post-processing. The post processing method is explained in §3.4.1.

Table 7 Average classification accuracy using the window method with changing spectral resolution

| | Spectral resolution | | | | | | |
	6.0 nm	8.3 nm	10.7 nm	13.1 nm	15.5 nm	17.9 nm	20.2 nm
Non-sampled training data	93.19%	93.40%	**93.58%**	93.25%	92.80%	92.48%	91.45%
Sampled training data	-	89.83%	**90.43%**	89.00%	89.70%	89.51%	88.17%
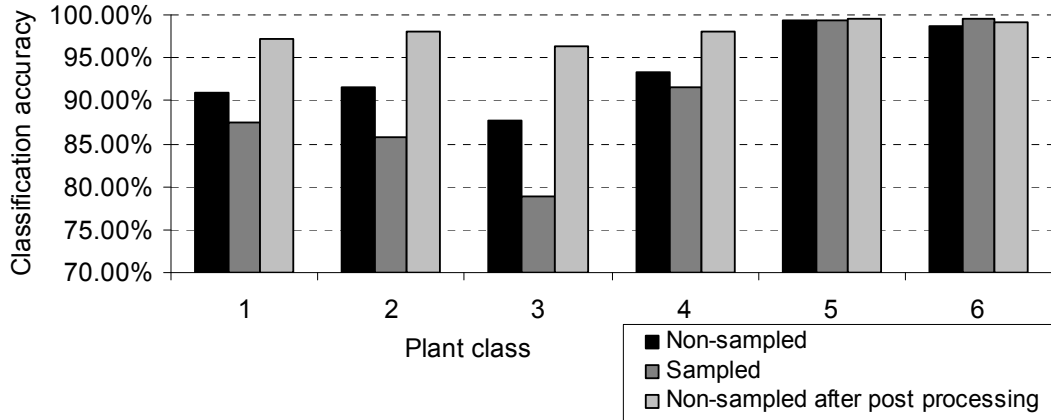
Figure 50 Classification accuracy in window method in different plant classes with 10.7 nm spectral resolution

Random sampling of data to create an equal number of training cases resulted in lower or equal classification performance in all classes but class 6. Class 6 had the fewest training cases in the non-sampled situation and it could be expected that its border-line measurements are easily classified to classes with a higher number of training cases. Overall, the results do not seem to justify using random sampling when more samples are available.

Spatial post-processing increases the classification probabilities in all classes. The

improvement was greater when original classification result was relatively low. The average classification performance after post-processing was 98.0%. The classification error matrix for the window method after post processing is shown in Table 8. The bottom row in the table shows classification results reflecting the number of complete leaves classified correctly such that all the windows within any leaf are classified correctly to the same class.

Table 8 Window classification matrix with a 10.7 nm spectral resolution

| | Class | Colour and spectrum (λ_{res}=10.7nm) | | | | | | Average |
		1	2	3	4	5	6	
Windows classified as (%)	1	**97.2**	1.8	2.5	0.5	0.0	0.9	
	2	0.3	**98.0**	1.0	0.9	0.0	0.0	
	3	1.1	0.1	**96.3**	0.5	0.0	0.0	
	4	1.5	0.0	0.2	**98.0**	0.3	0.0	
	5	0.0	0.1	0.0	0.0	**99.6**	0.0	
	6	0.1	0.0	0.0	0.0	0.1	**99.1**	**98.0**
Complete leaf classification (%)		**88.4**	**90.6**	**80.6**	**88.5**	**96.3**	**98.2**	**90.4**

5.3.3 CLASSIFICATION OF OVERLAPPING LEAVES

The results of the test with overlapping leaves are presented in Table 9. The same classification functions were used as with separate leaves, but only leaves from classes 1-4 were imaged in this test.

Table 9 Classification results of windows on overlapping leaves

| | Class | Classification of overlapping leaves (λ_{res}=10.7 nm) | | | | Average |
		1	2	3	4	
Windows classified after filtering as (%)	1	**86.9**	12.9	10.8	10.7	
	2	3.2	**82.8**	2.8	1.0	
	3	7.9	0.6	**85.6**	2.5	
	4	0.5	3.1	0.7	**85.7**	**85.2**
	5	0.0	0.6	0.0	0.0	
	6	1.5	0.0	0.1	0.0	
Before filtering (%)		**86.4**	**81.1**	**83.5**	**85.1**	**84.0**

The results are not as good as with separate leaves. There seems to be a strong bias to incorrectly classify windows to class 1. One reason for lower performance may come from the fact that leaves are somewhat transparent and spectral data is of mixed plant material at certain locations. The current detection system can make sure no soil is mixed into spectral measurements but it cannot necessarily separate different plants from each other.

5.3.4 LEAF BORDER ESTIMATION

An example of leaf border estimation is shown in Figure 51, where the centre locations of the classified windows are shown on a conventional image of leaves in (a), the estimated results in (b) and the correct borders in (c). The correct borders were derived manually from the segmented binary image of plants and soil. In this example few windows are misclassified, resulting in small errors compared to the true borders.
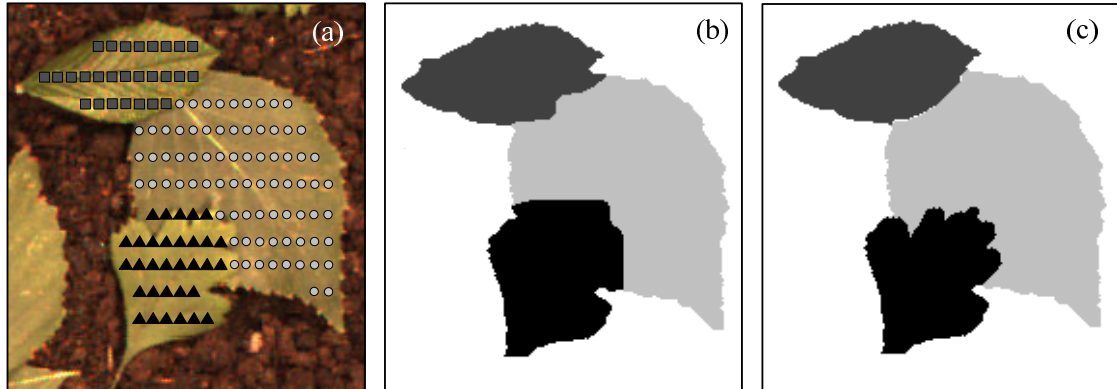
Figure 51 An example of leaf border estimation. (a) An RGB image of leaves showing centre points of classified windows, ■Rugosa Rose, ◎Silver Birch and ▲Hawthorn, (b) estimated leaf borders and (c) true leaf borders based on segmented images

The results of leaf area estimation with overlapping leaves are presented in Table 10. It is not valuable to compare the leaf border estimation results directly to the window classification results. Comparison can be made to a previous relatively complex method by Johansson and Baerveldt (2005), where 80% of the leaf area was correctly estimated. The average result in this project was 84.3%, which is slightly higher than the previous results.

Table 10 Estimation results of borders and area with overlapping leaves

| | | Estimation of leaf borders and area (λ_{res}=10.7 nm) | | | | |
	Class	1	2	3	4	Average
After filtering (%)		86.4	80.8	85.1	84.9	84.3
Before filtering (%)		85.7	79.7	83.0	84.3	83.2

A valuable use for the leaf border estimates would be to further process the border estimates such that partial shape recognition could be applied. Thorough investigation of the benefits of such use of data is beyond the scope of this project.

5.3.5 FINDING IDEAL DESCRIPTORS FOR PLANT CLASSIFICATION

It is important to know which descriptors have the highest significance for classification purposes. Calculation of only the most significant descriptors can save

processing power and time. The significance of a descriptor is dependent on the properties of the objects being investigated and the differences it is able to describe between classes.

A simple test of frequency of occurrence for each of the descriptors used in this study was done with both the full leaf and window methods. The test with full leaf method was done using all five descriptor sets and nine different spectral resolutions. Figure 52 shows an example of descriptor occurrence for all descriptors and their order of significance in the final classification model. Each row in the figure represents a model with a different spectral resolution data. The columns on the left side of the figure show the twelve shape and colour descriptors and the right side shows the spectral descriptors. In the figure, the shape and colour descriptors fill an entire square and spectral descriptors fill a width proportionate to the spectral band width. Darker colours, as shown in the order of occurrence legend, represent descriptors with higher order of occurrence and thus higher significance to the model.
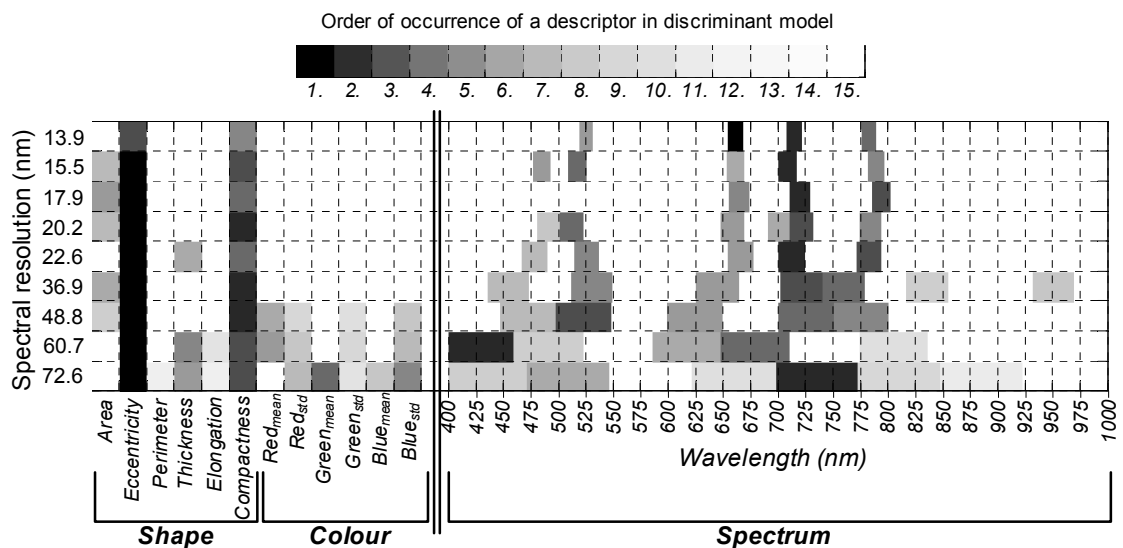
Figure 52 An example of descriptor occurrence in classification model with all descriptors and varying spectral resolutions using full leaf method. Darker colours represent a descriptor with higher significance in the classification model.

It is clear in the figure that the two shape descriptors of eccentricity and compactness are important regardless of the spectral resolution used. Colour descriptors become significant only when the spectral resolution is at a fairly broad band corresponding better to the wide spectral band from red, green and blue sensor elements.

By definition, the eccentricity and elongation descriptors seem to be describing similar shape properties of an object. However, only eccentricity is shown to be significant for classification purposes. The low significance of the elongation descriptor cannot be

contributed to the stepwise selection process of the variables, where prior selection of the eccentricity descriptor could have excluded a similar descriptor (elongation) from entering the model. When eccentricity was removed from the model, elongation did not take its place as a significant descriptor. Correlation coefficient, a measure of similarity of variables, between elongation and eccentricity is only 0.55. Explanations to low correlation and the low significance of the elongation descriptor compared to the eccentricity were not found.

Four spectral data regions can be separated with a relatively high occurrence in classification models. These regions or bands are approximately at 500-525, 650-675, 700-725 and 775-800 nm.

The region around the 'red edge' just beyond 700 nm as shown in Figure 8 seems to hold the most important spectral descriptors regardless of the spectral resolution. This is no surprise, since the reflectance signature of a living plant is greatly shaped by the chlorophyll content and its reflectance and absorbance properties. The slightest wavelength shifts in high gradient curve of the red edge region, caused by the chlorophyll concentration and the physical structure of the plant, can create relatively high differences in descriptor magnitudes. It seems that a spectral descriptor encasing the 710 nm wavelength right on the 'red edge' is present in the classification model at nearly all spectral resolutions with a high significance. The apparent variation in the descriptor location at the red edge at different resolutions is most probably due to the discrete nature of the descriptor locations.

The descriptor goodness can also be visualised like in Figure 53, where the best performing descriptor pair in a) and the worst performing descriptor pair in b) are shown with the training set data using the 22.6 nm data. The figure a) shows how most classes are separated using only two descriptors and b) how most classes are overlapping with less descriptive descriptors. It is important to note that in this example class 6 is not separated in a) but can clearly be separated in b). This illustrates the power of discriminant analysis with multiple descriptors or variables, where a class is often separated from others only with a combination of variables.
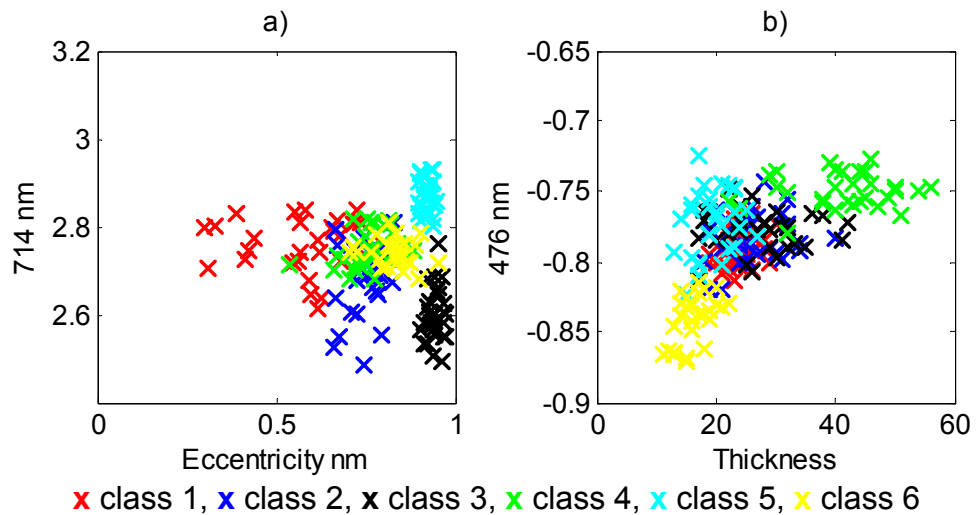
Figure 53 An example of differences between (a) best and (b) worst performing descriptor pairs with 22.6 nm spectral data from full leaves

Descriptor occurrence frequencies were also tested for the window method. The available descriptors were six colour descriptors and spectral descriptors in six different spectral resolutions from 8.3 to 20.2 nm. The standard deviation of the blue channel was the most significant of the colour descriptors, followed closely by green mean and green standard deviation.

A closer look was taken of the spectral descriptor occurrence frequencies for a comparison with previous research projects. The frequencies for all wavelengths were added together from the four descriptor sets using spectral data in both full leaf and window methods, and the relative occurrences are shown in Figure 54. The figure also shows a comparative spectral data occurrence from Thenkabail, Enclona et al. (2004).
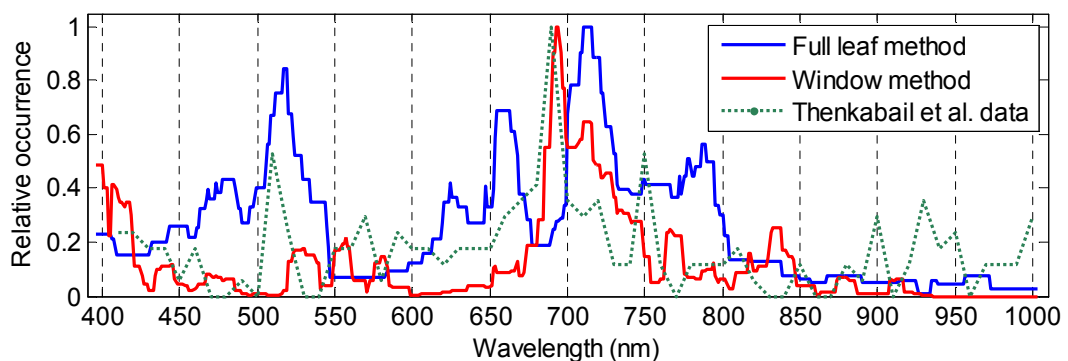
Figure 54 Measured spectral band occurrences compared to similar data from (Thenkabail, Enclona et al. 2004)

The full leaf data shows high occurrence values at the previously mentioned spectral bands as expected. There is a high correlation between the ~690 nm peak in the window data and the comparison data by Thenkabail, Enclona et al. (2004), but the peaks between

the comparison data and the full leaf data look to be offset. This offset may be a result of differences or errors in calibration or plant types, but no conclusions can be made. The overall trend does look to be similar, with important wavelengths at around 500-550 nm and 650-800 nm. The differences above 900 nm can possibly be explained by the poor sensitivity of the used CCD sensor during measurements in that range, whereas the equipment used by Thenkabail, Enclona et al. (2004) was a high quality spectrometer for the spectral range of 400-2500 nm. Their comparison data was recorded with a 10 nm spectral resolution. Vrindts, Baerdemaeker et al. (2002) have also identified the most important wavelengths in their research. Wavelength regions around 755 nm, 710 nm, 805 nm and 520 nm in this order were determined to be the most significant in the total range of 480 – 820 nm. These observations coincide also with the measurements in this project.

Comparison of colour and shape descriptor occurrences with previous studies show that compactness has been found to be important also by Åstrand and Baerveldt (2002). They also note that green mean and elongation are relatively significant. It has to be remembered that spectral data and eccentricity descriptors were not used in their study. Green mean has a high significance in this study as well. Apart from elongation, it can be said that the significant shape and colour descriptors in this study are similar to the ones in Åstrand and Baerveldt (2002).

5.4 DISCUSSION

The results from laboratory testing were equal to or better compared to previous studies on the subject. These results give evidence that the proposed system works well in its intended application under controlled circumstances. The classification probabilities of 85% or higher depending on the method used are high enough for significant reductions in chemical use in weed control.

Using information from both colour and spectral cameras has been shown to perform best. Low cost spectral solutions are limited to using silicon sensor cameras with a relatively low bit depth. The silicon sensor limits the spectral range to approximately 400-1000 nm. With a low cost system like this it cannot be expected that using spectral data alone would perform better than classification with data from a high quality spectral camera like presented in Thenkabail, Enclona et al. (2004).

The spatially accurate colour data is invaluable in detecting the leaf borders at a sub-millimetre level, while spectral data allows for robust classification. The colour data can also easily be used to verify the spectral data is collected from only plant material. This eliminates the need for complex algorithms to be used for spectral unmixing.

Shape has shown to be a valuable addition for classification. Simple shape calculations work well when plants or leaves are separate, but extracting shape information in cases when plants are overlapping is much more complicated and unreliable. The window method gives tools for recognizing overlapping plants with a good accuracy, but extraction of full shapes from occluded leaves is difficult. High spatial resolution of a colour camera is needed for accurate shape estimation.

Many previous research projects have recorded their plant classification accuracies for a two class case i.e. a crop plant and a class for all weed plants. When data from this project was investigated such that one class was selected as a crop plant and others as weeds, the classification with linear discriminant functions did not work well. The poor performance is most likely due to restrictions in the linear method. Typically the measurements for any class form a cluster of data in the multidimensional space of descriptors. When these clusters are well apart, a linear separation can be created around them. When a separation of one cluster from all other scattered clusters is attempted, the creation of reliable linear boundaries becomes difficult.

The performance was much better when each class was classified in their own class and the two-class conversion done after this classification. Such performance figures can be estimated from the classification error matrices. Classification accuracy for the selected crop class is the same as the classification shown in the matrix diagonal. The misclassification of the crop can be calculated by adding together the other values in the selected column. The other plants then form the weeds-class. Their combined classification accuracy can be estimated by adding together all the values in remaining classes excluding the values at the row of the selected crop plant and dividing the sum with the number of classes to be combined. An example result from Table 8 choosing the class 1 as the crop and others as weeds is shown in Table 11.

Table 11 Two-class classification example

	Crop	Weeds
Crop	**97.2**	1.1
Weeds	3.0	**98.8**

Certain results from the laboratory testing can be used to re-evaluate the design of the system towards faster and more economical solutions. First, the optimal spectral resolution was found to be lower than the nominal spectral resolution of the spectrograph and second only certain spectral regions were shown to be significant for classification. Lower spectral resolution requirements can be fulfilled with lower resolution monochrome cameras thus

lowering the cost. Also, random access of spectral image information to output only the most important wavelengths can increase the speed of the system.

The performance differences between different spectral resolutions in full leaf method were not significant. This may lead to the conclusion that using only lower spectral resolution data cannot be overlooked as an option. This reduction in spectral resolution could potentially save greatly in image transfer and processing time. There exists no previous study on spectral resolution sensitivity analysis in plant classification.

The optimal spectral resolution using the window method was less than half of the nominal spectral resolution of the spectrometer. It needs to be noted that the spectral resolution of 10.7 nm is optimal only for this particular camera setup. Higher or more accurate spectral resolutions might be optimal when cameras with higher bit-rates or lower noise levels would be used.

The laboratory tests have given a sound base for outdoor tests in real field environment. Outdoor tests are needed to provide verification of the system capabilities in uncontrolled conditions.

6.0 PLANT CLASSIFICATION IN FIELD CONDITIONS

Laboratory tests had proven that the system was capable of accurate plant detection and classification. Outdoor field tests were the natural next step to verify system usefulness in a more realistic environment. Outdoor conditions expose the system to several additional variables compared to the controlled laboratory environment. These include the changes in lighting due to weather, season and time of day, rain drops on leaves and random plant patterns.

This chapter gives an overview of the field test arrangements and the modifications made to the test rig and the data processing methods. A thorough analysis of the field test results is also presented and the results compared with previous studies.

6.1 FIELD TESTS

6.1.1 TEST RIG MODIFICATIONS FOR FIELD TESTS

The test rig was originally designed for laboratory use where the rig need not be moved. Mobility became important once the rig was in the field. The main change to enhance mobility was the addition of four wheels to the rig. The wheel base was fixed but the wheels could be individually adjusted along their axis to fit between plant beds in order to avoid running over plants.

A hard cylinder with a soft curtain extending the ground level blocking the outside light from the imaging area was also added. This was done as the original curtain from a soft fabric could not maintain its shape in the wind and occasionally let some unwanted light in. The soft curtain dragging on the ground allowed plants to enter and exit the imaging area unharmed, and also provided added cover from the outside light. It also dealt with slightly uneven ground better than a rigid structure.

A shelf was also added to the rig to hold the computer. This was covered in clear acrylic plates to protect the equipment from light rain. During the outdoor tests the rig was powered by a petrol generator. All the main modifications are shown and labelled in Figure 55. An image of the test rig during actual field tests on a sugar beet bed can be seen in Figure 56.

A small modification was also done to the microcontroller program and the communication between the computer and the microcontroller. An RS-232 link was added between them to control image acquisition and movement of the cameras.
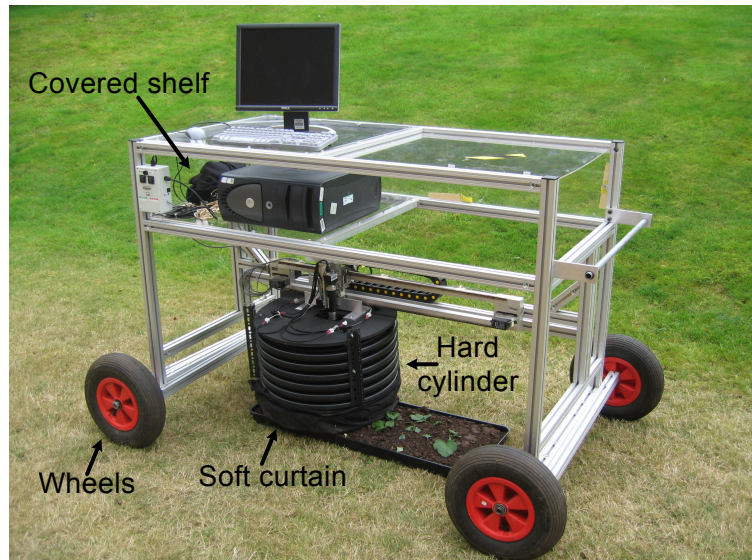
Figure 55 Test rig modifications for field tests

Figure 56 Test rig on the field

6.1.2 DATA COLLECTION GUI

A simple graphical user interface (GUI) was created in Visual Basic to help with the collection of the data. The interface was designed to handle the communication between the microcontroller and the computer and allow simple buttons to perform image acquisition, changes in camera properties and manual and automatic movement of the camera head.

A screen shot of the GUI is shown in Figure 57. Separate functional areas are labelled in the image. Buttons in area 1 control the serial communication between the microcontroller and the computer. The coloured square shows if the communication link is open (green) or closed (red). The buttons in the area 2 control the camera head movement. Coloured circles indicate which activity is currently in operation. The Go Home-button will move the camera head to the home position at the beginning of the

linear guide, the Acquire-button will start the automatic image pair recording and the STOP-button will stop all movement and recording. The button in area 3 is used to send a break signal into the serial bus in case the system gets jammed due to a communication error. The buttons in area 4 are used to move the camera head step by step and buttons in area 5 to record individual image pairs. Buttons 1-5 automatically name the saved image files with respective numbers. Area 6 is used to choose the folder in which the images are currently saved. Areas 7 and 8 are showing the live images from the colour and spectral cameras. The buttons and text fields in these areas can be used to access and alter the respective camera settings.
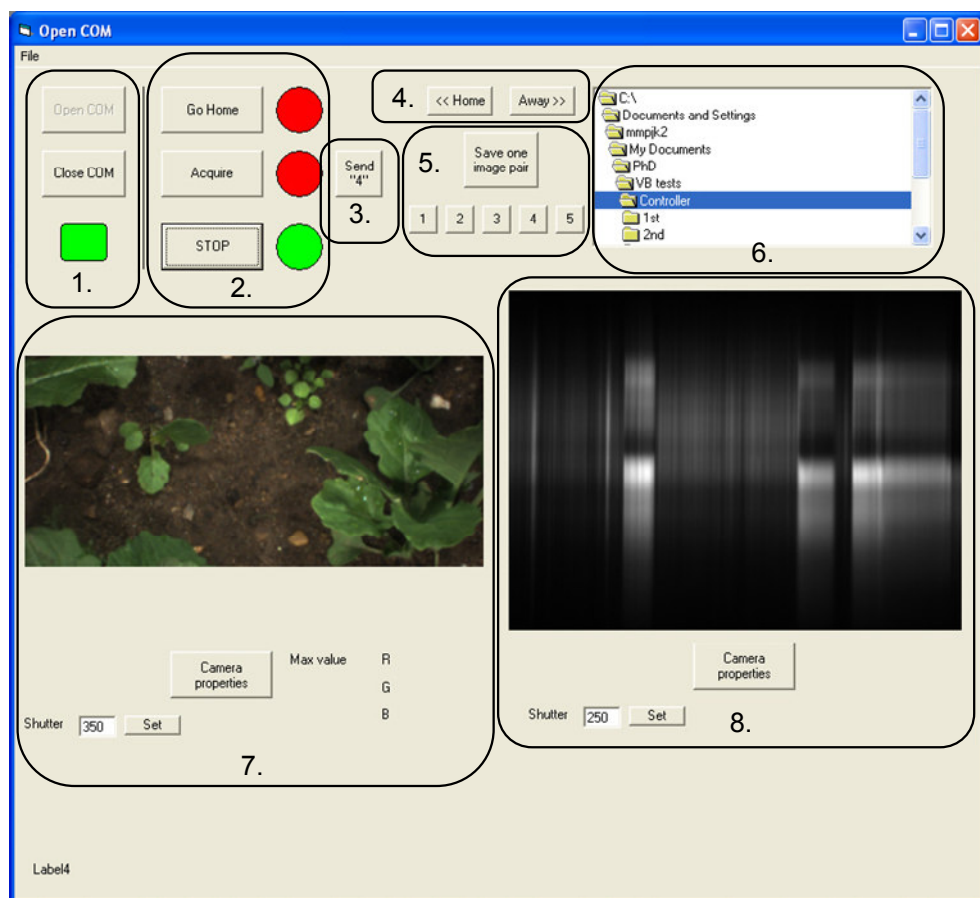
Figure 57 Screenshot of the GUI for measurements

6.1.3 CALIBRATION PROCEDURES DURING FIELD TESTS

Lighting calibration was considered the only calibration procedure that needed to be done every day before taking the measurements. During transportation of the rig, the relatively loosely fixed lights could change their position slightly. Lighting calibration images for the colour camera were taken as during the laboratory tests by acquiring images of white paper filling the entire camera field of view (see §5.1.6). Spectral images were acquired at the same time and the white paper images were used in spectral lighting

calibration as well instead of using the Spectralon reflectance tile. Images taken of white paper gave full calibration images without the need to manually move the reflectance tile in the camera field of view.

The Spectralon tile specifications claim the tile has a uniform diffuse reflectance of 90% of the light over the spectral band of 400 – 1000 nm. Figure 58 shows a comparison of measured white paper reflectance to Spectralon tile reflectance over the used spectral range. The white paper has a fairly uniform reflectance of approximately 75% of the light. Differences in the reflectance values do not affect the overall results as long as all the lighting calibration is always done with the same method. The comparison figure just shows that even white paper produces fairly uniform reflectance over a wide spectral range.
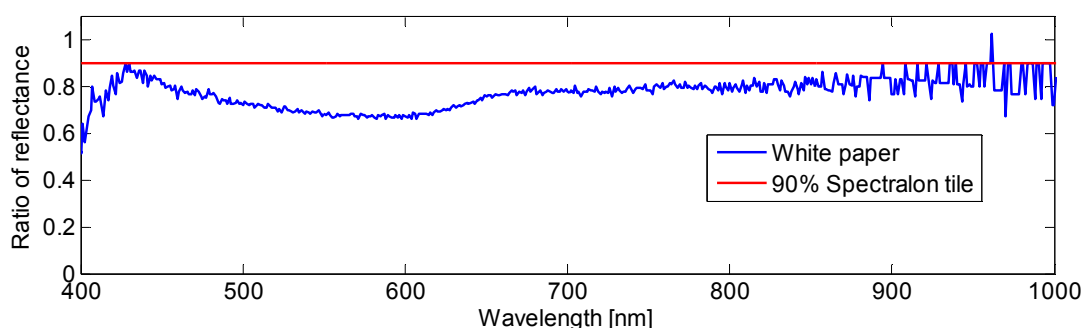
Figure 58 Spectralon vs. white paper reflectance

6.1.4 THE FIELD TEST SETUP

The field tests took place in an allotment plot in Loughborough (United Kingdom) over June and July 2007. Seeds from four crop plants and five weed plants were sown on the 21st June, and a list of the sown plant seeds is shown in Table 12.

Table 12 Plants sown in for field tests

Sown Crop Plants	Sown Weed Plants
Winter wheat (*Triticum aestivum*)	Charlock (*Sinapsis arvensis*)
Winter oat (*Avena sativa*)	Cleaver (*Galium aparine*)
Oilseed rape (*Brassica napus)*	Chickweed (*Stellaria media*)
Sugar beet (*Beta vulgaris*)	Mayweed (*Matricaria inodora*)
	Black grass (*Alopecurus myosuroides*)

Wheat, oat and oilseed rape were sown into 6 rows with 12.5 cm row spacing. Wheat and oat seeds were sown 30 mm and oilseed rape seeds 60 mm apart. Sugar beet seeds were sown with 35 cm row and seed spacing into four rows. An image of the whole test field taken at the end of the trials is shown in Figure 59. In the image all the plant rows are in near vertical direction.

| Oat | Wheat | Oilseed rape | Sugar beet | Weeds |

Figure 59 Trial field and crop plant bed locations imaged 29 days after sowing

Figure 60 illustrates how weed seeds were sown into eleven regions (A-K) in each crop bed and also how imaging areas were located and named for each bed. The weed regions were identical among the four crop plant beds, e.g.. the top left corner of each crop bed had a weed region A and an imaging area '11'.

In the figure, different colours represent the weed regions. One weed region was left free of sown weed seeds and others had a selection of weed seeds as explained in Figure 60. Sowing pattern of weeds was random within a region. This was done to realistically represent the occurrence of weeds in fields. A separate weed plant bed was also sown as a control area where only the sown weed plants would grow with other naturally occurring weeds. This weed bed consisted of weed regions B-F.

Each plant bed was also divided into an imaging grid indicated by column and row structures in Figure 60. Each imaging cell represented an area where images were captured during one measurement without moving the test rig. During the measurements the test rig was moved manually over each imaging area. In practice, the imaged areas were slightly overlapping. In the oat, wheat and oilseed rape beds an imaging column enclosed two crop rows and in the sugar beet bed only one crop row. In fact, the sugar beet bed enclosed four imaging columns and the weed bed only one column with three imaging rows.
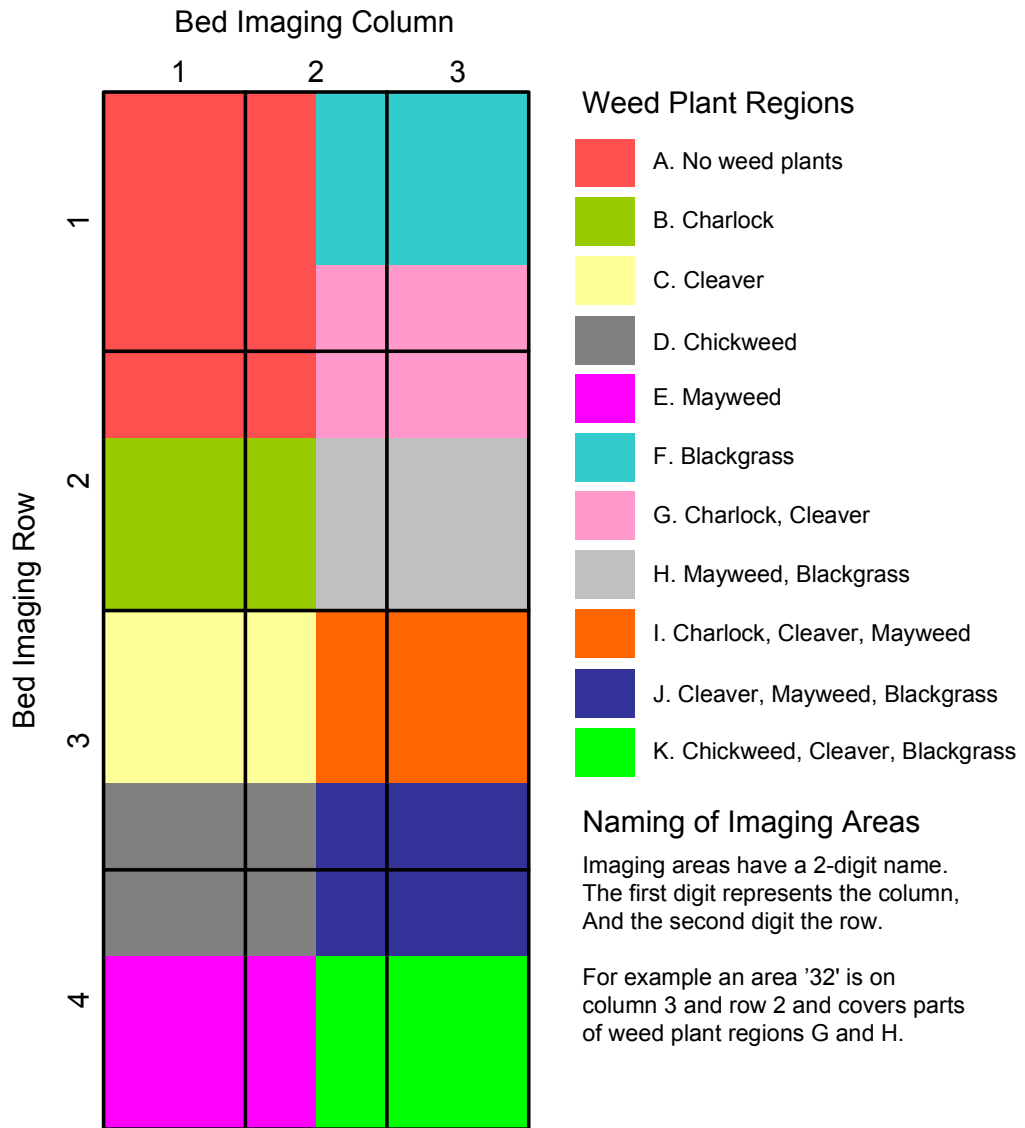
Figure 60 The regions for sown weed seeds and the imaging areas within each crop plant bed.

The soil naturally had plenty of commonly occurring weed seeds within it. During the outdoor measurements eleven additional weed species were recognized. These are listed in Table 13.

Table 13 The recognized additional weeds

Weed Plants in Soil
Fat-Hen (*Chenopodium album*)
Docks (*Rumes spp.*)
Perennial Sowthistle (*Sonchus arvensis*)
Annual Sowthistles (*Sonchus spp.*)
Crane's-Bill (*Geranuim spp.*)
Shepherd's Purse (*Capsella bursa-pastoris*)
Red Dead-Nettle (*Lamium purpureum*)
Common Couch (*Elymus repens*)
Henbit Dead-Nettle (*Lamium amplexicaule*)
Small Nettle (*Urtica urens*)

Measurements were taken on five different days over a period of 2 weeks. These were 11, 14, 18, 21 and 25 days after sowing (DAS). All imaging areas were used on the first four days, and locations 11-21 were excluded on the fifth. The exclusions were due to plants growing too high for the current system to image them reliably.

6.1.5 APPLICATION SPECIFIC IMAGING PROBLEMS

6.1.5.1 *Perspective Error from Camera Positioning*

The laboratory measurements were taken of plant leaves on a flat surface. In this setup, the location of the cameras relative to each other and their precise angle with respect to the leaves was not meaningful as long as the position of the spectral line could be located precisely within each colour image. However, the separation of the camera optical axes becomes meaningful when the objects are at varying distances from the cameras. In the case of the presented system, the cameras stay at a constant distance from the ground while the height of the plants varies.

When the camera optical axes are separated, the same image cannot always be seen at both cameras at the same time. Occlusions, as presented in Figure 61.a, where the colour camera cannot see point x when the spectral camera can, become then possible. When the cameras are moved along the linear guide, the original view of the spectral camera becomes eventually available for the colour camera as shown in Figure 61.b. In the test rig setup this offset was 20 image pairs, i.e. objects in spectral camera field of view in image $I\_spectral_1$ appeared at the optical axis of the colour camera in the image I_colour_{21}.
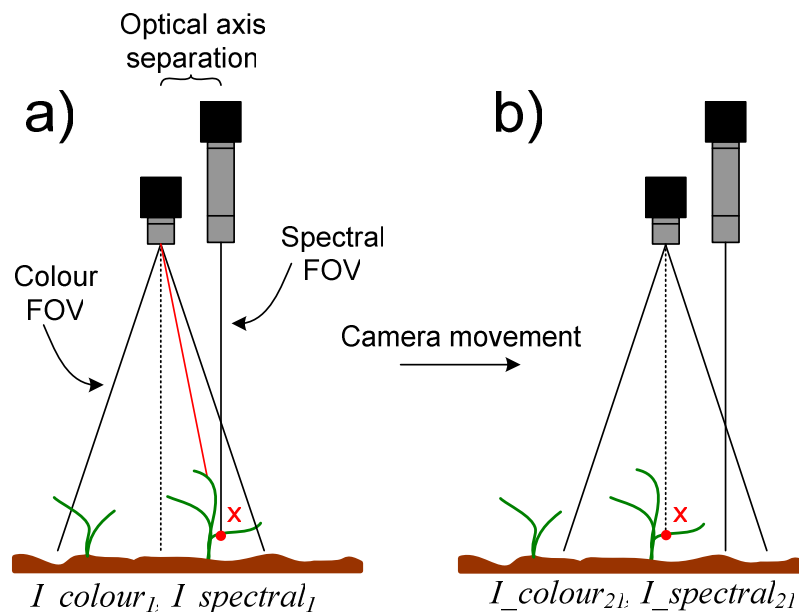
Figure 61 Occlusion and perspective error due to separated optical axes

The need for the image offset could have been avoided by using a half mirror setup as

shown in Figure 62, where the optical axes of the cameras would be aligned. The disadvantage of this setup is that the amount of light for both cameras is halved. This was the reason such a setup was not initially chosen for the application where available light intensity was thought to be limited and exposure times critical.
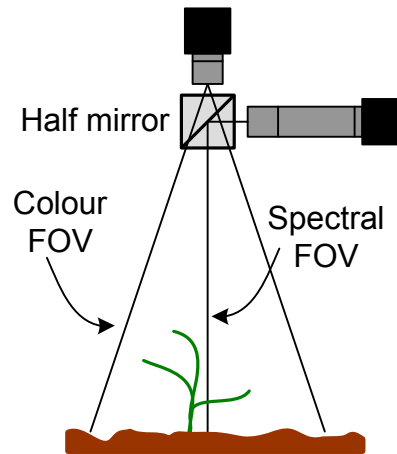
Figure 62 An alternative camera setup using a half mirror to align the optical axes

The camera movement can be used to create a stereo image pair and give estimates of plant height from colour images. A stereo camera pair can be created with two stationary cameras or one moving camera. As long as an estimation of the camera movement within a pixel or two can be given, the height of the objects in view can be calculated. Further analysis of this stereo problem is beyond the scope of this project.

6.1.5.2   *Effect of Moving Plants*

During the tests, tall and thin plants were moving in the wind. The curtain blocking the outside light was also dragging close to the ground making the plants move. These effects could be clearly seen in the registered images. Figure 63 shows an example of two consecutively captured colour images and a registered image consisting of these two and many more images. The black oval in all three images is covering the same location relative to the ground. In the two example images, the wheat leaves are clearly in different places. This movement of leaves between frames creates discontinuous plant parts in the registered image, which is constructed a few image rows at a time. A data window could not always be reliably placed on locations where the leaves are moving since the registered image was used for plant segmentation.
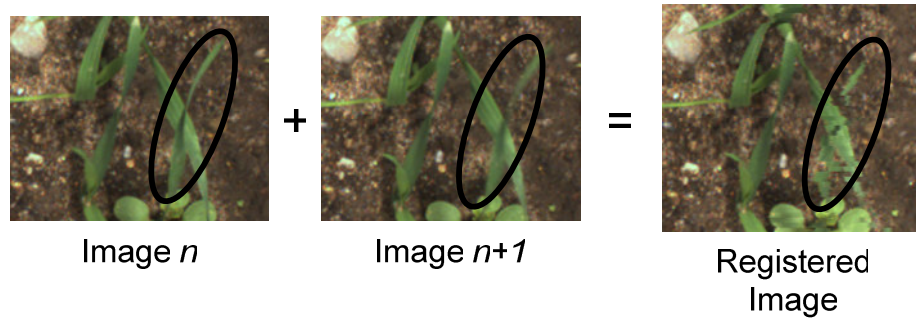
Figure 63 Effect of moving leaves in registered images

The moving leaves occurred rarely and almost solely with grassleaved plants. This effect was therefore not considered important for further analysis. Test rig modifications were also not considered.

6.2 DATA PROCESSING

The following section explains the post-processing used for field test data where it differs from the procedures during the laboratory tests. Whenever possible, Figure 19 is referred to when additions or changes are explained.

6.2.1 IMAGE PRE-PROCESSING

Image pairs at maximum camera resolutions, 24-bit 659 x 493 pixel colour and 8-bit 1360 x 1024 pixel monochrome images, take a lot of storage space. It was decided that only the first 300 rows of the colour image would be recorded and the image resolution of the monochrome images was reduced to 680 x 512 during the outdoor tests. This effectively reduced the required storage of any image pair significantly without losing any important information. Even fewer colour image lines were needed for the final processing, but the CCD camera did not allow a complete choice over the region of interest.

After each measurement session a further reduction was done to the images. Only colour image rows 140 – 300 were saved. This selection was based on the need to save the original spectral line location and the central line (optical axis) of the colour image with a buffer of some 20 lines on both sides. Every nine spectral points along the vertical spectral axis of the monochrome images were averaged together. This reduction of 512 data points to 56 along the spectral axis gave a spectral resolution of 10.7 nm. This spectral resolution was determined to be optimal for the system during laboratory testing. The combination of these reductions brought the required storage space for an image pair from 2.3 MB to 620 kB. In Figure 19 these steps happen at 2.a in spectral images and 1.b in colour images.

6.2.2 LIGHTING NORMALISATION IN COLOUR AND SPECTRAL IMAGES

During the laboratory tests it was easy to adjust the light setup to give the most uniform lighting throughout the field of view. The setup was not subject to vibration during transportation and thus remained constant for all laboratory tests. This was not the case for the field tests. The lighting setup changed significantly from day to day, but did remain constant for the duration of the measurements in one day. The changes to the normalisation procedure explained here are applicable to 1.a and 1.b in Figure 19.

The basic effect of these changes in lighting was introduction of significant intensity changes in the images. It was quickly realized that the lighting normalisation explained in §5.1.6 (the old method) for each of the RGB channels separately did not maintain the colour information. A new method was needed. After trial and error, a modified calculation of chromaticity values was chosen as described with the following equation:

$$\begin{bmatrix} R_{Norm} \\ G_{Norm} \\ B_{Norm} \end{bmatrix} = \begin{bmatrix} R / \sqrt{(R_C^{\,2} + G_C^{\,2} + B_C^{\,2})} \\ G / \sqrt{(R_C^{\,2} + G_C^{\,2} + B_C^{\,2})} \\ B / \sqrt{(R_C^{\,2} + G_C^{\,2} + B_C^{\,2})} \end{bmatrix}, \tag{6.1}$$

where R_C, G_C and B_C are normalised calibration image channels (taken of white paper) containing values from 0 to 1. This effectively maintains the colour relationship while normalising the intensity changes described by the calibration image channels. The denominator for each channel represents the length of an RGB vector at a current location. This modified chromaticity calculation was observed to perform better than the original calculation where the denominator would simply consists of the sum of all three channels. A comparison of the old method and the new method is shown in Figure 64.

The colour information is not preserved with the old method even though the overall intensity changes are removed. The colours remain correct only in the central portion of the image where the intensity is greatest. The image after the new normalisation method has a uniform colouring without intensity changes. Although difficult to see, the difference between the methods is significant.

a) Original Image

b) Calibration image RGB vector length

Colour information lost

c) Normalized using the old method

Colour information preserved

d) Normalized using the new method
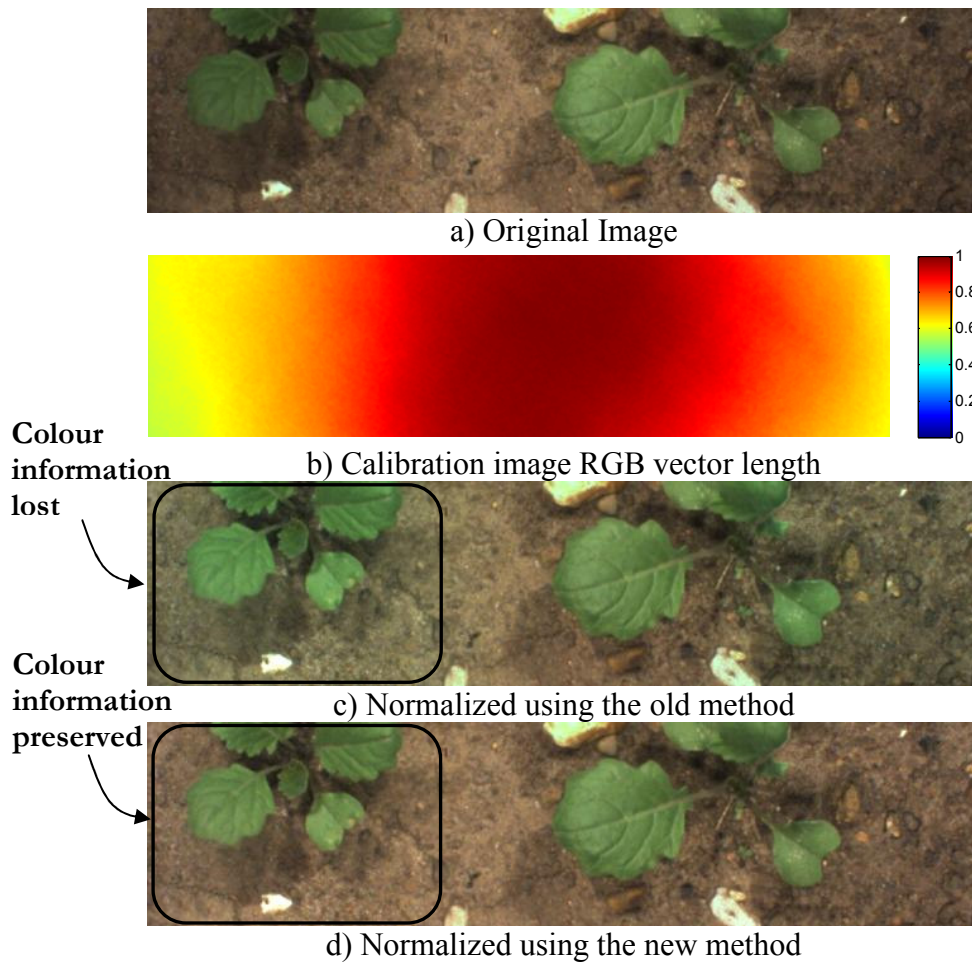
Figure 64 Comparison of the light intensity normalisation. a) The original image with intensity changes, b) calibration image RGB vector length (the modified chromaticity denominator), c) the image normalised using the old method and d) the image normalised using the new method

During the laboratory testing the lighting in spectral images was normalised according to the highest values at each spatial location. The method did not correctly account for the possible changes in individual wavelengths, but multiplied the whole spectrum at a given spatial location with one value. This resulted in problems with classification especially in the edges of the image due to a similar effect that can be seen in Figure 64 with the old colour normalisation method. Data especially in the edges of the spectral image did not preserve the correct spectral curve shape after a simple normalisation. Therefore, a method of normalising each wavelength or spectral data point separately needed to be used.

The new method used full spectral calibration images consisting of values from 0 to 1. The acquired spectral image pixels are simply divided with corresponding calibration image pixels. As a result, the spectrum in the image follows the spectral properties of the lighting at each spatial location.

6.2.3 CHECKING FOR NON-MIXED SPECTRAL DATA

The plant segmentation was noticed to be inadequately accurate for making sure the windows were completely on plants when data from the field was analysed. The segmentation method shown in §3.1.3 was able to correctly extract the entire plant area, but shadowed areas and small gaps between the plants were often segmented as plants as well. An additional verification step was required just prior to classifying the data, i.e. before point 6 in Figure 19, to verify that the spectra in the windows was of plant material only and did not include soil or other material in the signal. This was done by checking the existence of the 'red edge', the rise in reflectance in the near infra-red region caused by chlorophyll content. This spectral signature feature is unique to all living plants.

To check the red edge, the difference between the normalised reflectance values of spectral data points at 748 nm and 683 nm was first measured. If this difference exceeded the manually set limit of 1.75 the spectral signal was accepted to be of plant material only. The threshold value of 1.75 was manually chosen after examination of all different plant types used in the test allowing for adequate variation in spectra. An example case in Figure 65 shows oilseed rape spectra signals and the wavelengths used for mixed signal detection. An arrow with a length of 1.75 also indicates the scale of the required difference in value for the non-mixed signal.
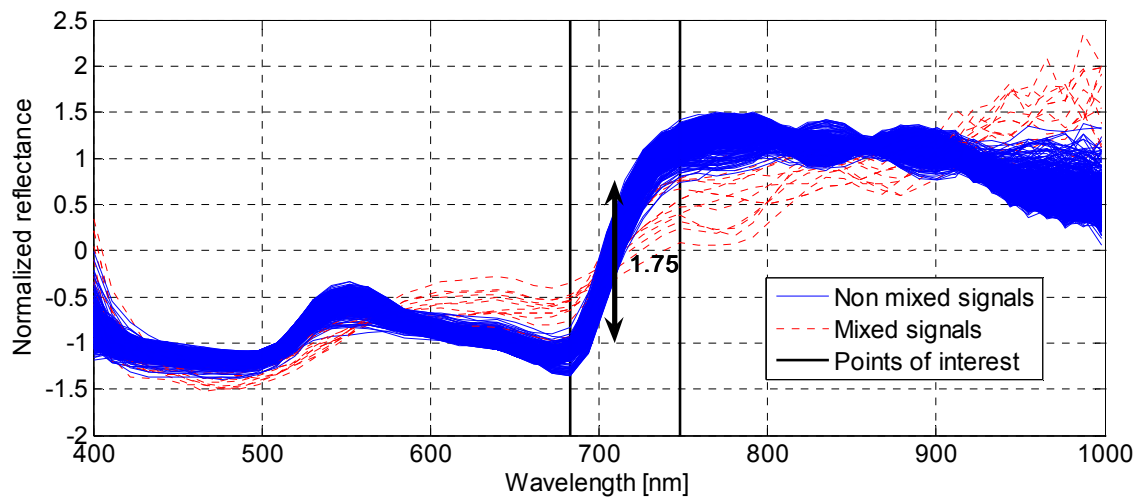
Figure 65 Example plot showing non-mixed and mixed rape spectra and wavelengths that are used for the detection of mixed signals

6.2.4 MANUAL PLANT RECOGNITNION AND LABELLING

Correct performance assessment of any classification system requires knowledge of the true class of each sample. Therefore, the real class for each sample, called the ground truth data, needed to be extracted. In this project the ground truth data could only be established by manual labelling of plants. Table 12 and Table 13 show the weed and crop

plant types present in the field. During the five measurement days nearly 300 measurement sets were gathered consisting of images from one acquisition run of the test rig length. Plants in over half of these images were manually identified and labelled accordingly.

Figure 66 shows an example of the labelling of output images. The original image a) shows high weed infestation. The labelled plant types in b) were recognized in registered images and their corresponding areas in the segmented binary image were coloured giving each plant type a distinctive colour. Unknown plants were left white. The separation of individual plants in Figure 66 c) by labelling each separate plant with a different value was done for oilseed rape and sugar beet data only. This was done to gather statistics on system performance in detecting separate plants. Each individual plant is shown in a different colour. The background blue represents non-plant or unknown areas. Finally, d) shows an example of labelled windows. Each pixel in d) represents one window. The plant label in b) that has the majority at each 7 by 7 pixel window area location is selected as the window plant label in d). Since each pixel in d) represents a 7 by 7 pixel area in b), the image in d) looks somewhat 'pixelated'.

During the recognition and labelling process some plants could not be identified. Such plants were left blank and were not considered during classification. Also, perennial sowthistle and annual sowthistle were extremely difficult to correctly tell apart by a human from the available images. The manual recognition with these plants is mixed. Sowthistles were considered as separate groups when creating the classification functions, but are shown both separate and combined in the results. Overall, the manual plant recognition did not create an error free ground truth data, but the errors were considered negligible for reliable estimation of classification performance.

a) Original registered image
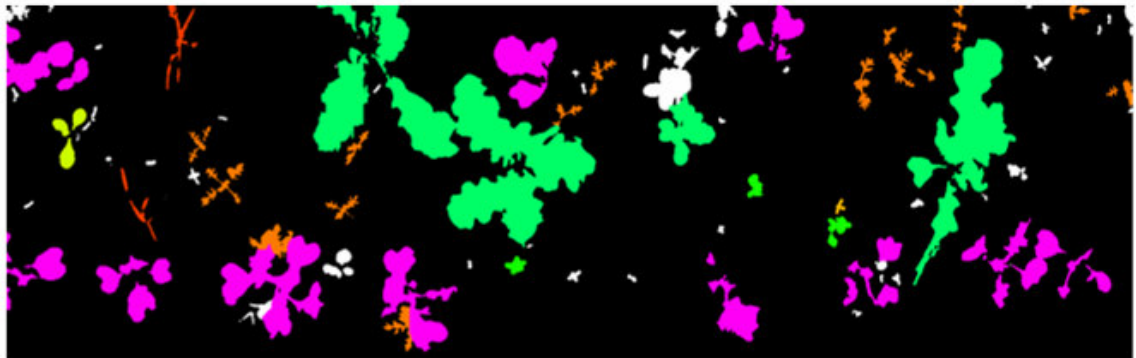
b) Labelled plant types
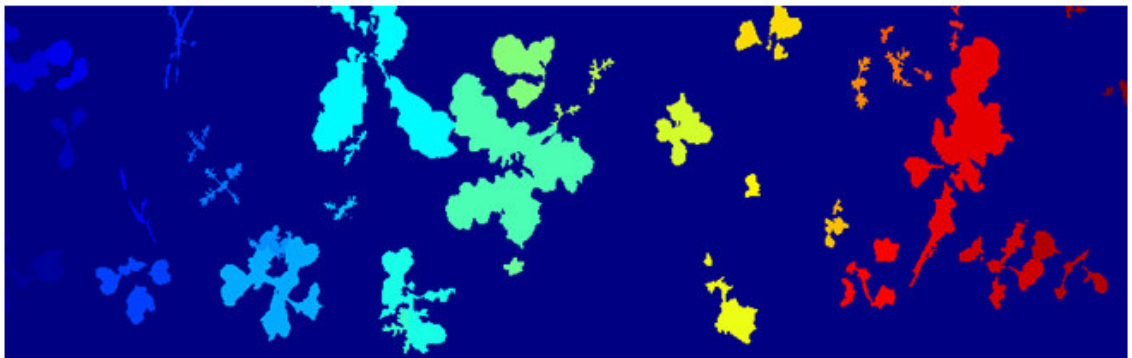
c) Separated individual plants
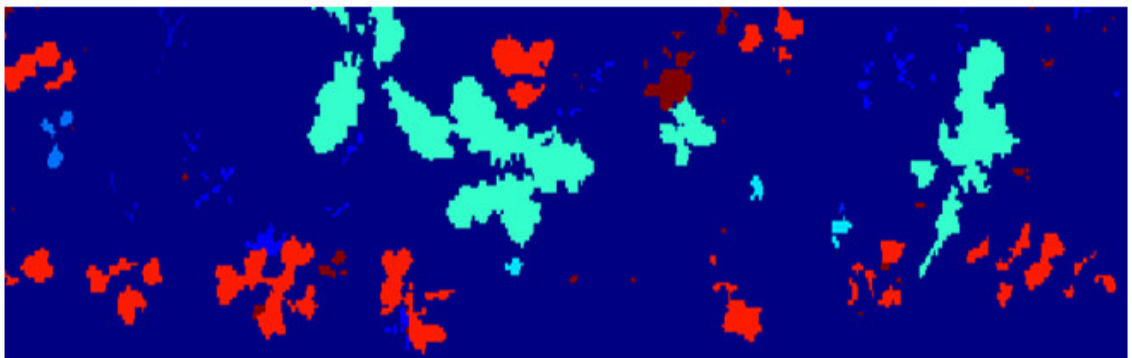
d) Labelled windows

Figure 66 Plant labelling example. a) Registered image of oilseed rape imaging area 33 taken 21 DAS, b) labelled plant types, c) separated individual plants and d) labelled windows

110

6.3 RESULTS

This section shows the classification results when the window method was applied to the outdoor measurements. The results also show the temporal changes in the reflectance spectra and the effects these changes have for classification. System performance is evaluated and thorough analysis conducted using measurements from both individual plant beds and all available data. Individual plant bed analysis provides insight into performance figures when one crop plant and some weed plants are present. The analysis using the complete data set shows the performance with multiple crop and weed plants. This analysis method using the full data was considered more descriptive for the total system performance evaluation. Majority of the analysis was therefore done using all data.

6.3.1 CLASSIFICATION IN A SINGLE CROP PLANT BED

The system performance was first evaluated by creating classification functions with training sets from individual plant beds where one crop plant and some weed plant species were available. The results for individual crop plant beds are shown in Table 14. The results show the crop plant, combined weed plant and average classification performance figures and the number of windows classified in each case. The number of weed species available for each case is also shown. The average value was calculated by averaging the individual plant class percentages.

The highest crop classification probabilities of 96.0% and 91.3% can be found from wheat at 11 DAS and sugar beet 25 DAS, respectively. The highest weed classification probabilities of 100% and 98.1% are from wheat at 14 DAS and sugar beet at 21 DAS. The highest average probabilities were 96% and 93.7% from wheat at 11 and 18 DAS. Overall, the results were in the same levels as the results from the laboratory tests. The presented investigation of results can also be compared with the results shown by Vrindts, Baerdemaeker et al. (2002). The results in this project are comparable or higher than the ones presented by them, where sugar beet and weeds were classified with probabilities up to 95% and 84%, respectively.

Results can be highly dependent on the selection of plants and the similarity in their reflectance properties. Oilseed rape classification results were significantly lower than with other crops. This was due to very similar spectra with charlock and resulted in high misclassification probabilities between the plants. The effect is explained in more detail with further results.

Table 14 Classification into crops and weeds. Average class performance, number of windows in the sample sets and the number of weed species available for each case are also shown.

	DAS / (Number of Classified Windows)									
Oat	**11**	**# win.**	**14**	**# win.**	**18**	**# win.**	**21**	**# win.**	**25**	**# win.**
Crop	**88.7**	(141)	**89.4**	(559)	**80.3**	(3736)	**88.9**	(9464)	**88.3**	(25249)
Weeds	-	(0)	-	(0)	-	(0)	**87.8**	(13)	**92.6**	(233)
Average	**88.7**		**89.4**		**80.3**		**89.0**		**82.8**	
Weed species		*2*		*4*		*3*		*3*		*5*
Wheat										
Crop	**96.0**	(859)	**90.5**	(2581)	**94.9**	(8619)	**86.6**	(16159)	**88.3**	(31088)
Weeds	-	(0)	**100**	(46)	**94.2**	(465)	**93.9**	(1167)	**95.8**	(5128)
Average	**96.0**		**90.9**		**93.7**		**75.8**		**72.8**	
Weed species		*2*		*4*		*7*		*8*		*9*
Oilseed rape										
Crop	**49.6**	(3084)	**88.5**	(5290)	**71.3**	(12760)	**65.1**	(30369)	**41.1**	(60747)
Weeds	**85.2**	(263)	**76.7**	(1010)	**78.6**	(3749)	**85.2**	(8767)	**91.5**	(17488)
Average	**65.7**		**82.4**		**78.7**		**78.9**		**60.2**	
Weed species		*6*		*12*		*13*		*13*		*11*
Sugar beet										
Crop	**81.8**	(258)	**83.3**	(837)	**88.9**	(3149)	**89.8**	(8550)	**91.3**	(25440)
Weeds	**93.2**	(322)	**91.1**	(1409)	**96.1**	(4206)	**98.1**	(12818)	**96.6**	(34843)
Average	**86.7**		**82.9**		**83.1**		**81.4**		**78.3**	
Weed species		*3*		*9*		*12*		*14*		*14*

The weed infestation in the oat and wheat imaging areas included in the analysis was relatively light. The classification training set included 120 samples for each plant type, and the rest, if available, were left for the sample set. The number of weed species samples with oat did not exceed 120 in any weed plant type before 21 DAS, where 13 windows from weed plants were available in the sample set. Lack of weed plant windows was likely to elevate the overall classification results for the grassleaved plants.

The general trend of performance figures is that the classification as early as possible resulted in the best results. This can be partially due to lower number of total plants available at earlier dates, but is nevertheless encouraging early action in weed control.

6.3.2 CLASSIFICATION OF WINDOWS USING ALL PLANT BEDS

The system performance figures with a single crop plant and some weed plants were high. However, it is important to analyse how well the system is capable of classifying plants when all known plant types or classes are included in the training set. This analysis gives an estimate on how flexible the system is to any number of plant classes. Table 15

shows the percentages of correctly classified windows for each day and all available plant combinations. An empty cell indicates that samples were not available for that particular class and day combination. Performance figures at or over 85% are in green and figures under 50% are in red for ease of reading. Two more weed plant species were detected than are included in this table. The two missing plant types had fewer than 120 data windows in total and these windows were therefore included in the training set only. The table also shows the total number of available windows in sample sets for each case. The available number of windows for each plant class increases with passing time simply because the growing plants are covering more area.

The average classification performance varied between 62% and 85% over the measurement days. This included average performance variations of 70-87% among the crop plants and 59-78% among the weed plants. The combined class of sowthistles (classes 6&10) was used in calculation of the averages instead of the individual classes 6 and 10. Data for 2^{nd} 25 DAS was measured at a spectral resolution of 8.3 nm. All other data sets used 10.7 nm data. The discussion on higher spectral resolution available at 25 DAS is given later in §6.3.9.

Table 15 Classification probabilities of windows in all plant types on different days. The total number of windows for each plant is shown in parenthesis. Performance figures of 85% or above are in green and under 50% in red. 2nd 25 DAS data spectral resolution was 8.3nm, all others 10.7nm.

Plant type	Classification performance % (Number of classified windows)					
	DAS					
	11	**14**	**18**	**21**	**25**	**2nd 25**
1. Fat-hen			55.4	64.6	17.6	48.4
(*Chenopodium album*)			(57)	(1108)	(4084)	(3383)
2. Blackgrass				68.3	61.1	66.2
(*Alopecurus myosuroides*)				(166)	(1941)	(1362)
3. Mayweed			67.4	80	69.2	75.5
(*Matricaria inodora*)			(224)	(1441)	(9005)	(5765)
4. Chickweed				77.2	75.9	82.4
(*Stellaria media*)				(92)	(1413)	(1051)
5. Docks				89.7	78.6	86.0
(*Rumes spp.*)				(301)	(2046)	(1763)
6. Perennial Sowthistle			66	53.9	48.5	60.2
(*Sonchus arvensis*)			(315)	(1154)	(4321)	(3763)
7. Cleaver				66.7	70	85.4
(*Galium aparine*)				(45)	(253)	(199)
8. Crane's-Bill		88.8	67.3	49.9	52.1	62.6
(*Geranuim spp.*)		(80)	(306)	(884)	(1843)	(1565)
9. Shepherd's Purse			79.3	59.5	50.8	67.1
(*Capsella bursa-pastoris*)			(213)	(972)	(2709)	(2202)
10. Annual Sowthistles			66.3	57.7	50.5	51.6
(*Sonchus spp.*)			(86)	(542)	(1320)	(1069)
11. Charlock	78.2	63.7	61.7	67.5	46.2	61.8
(*Sinapsis arvensis*)	(831)	(2884)	(9022)	(23057)	(51753)	(44332)
13. Common Couch			100			
(*Elymus repens*)			(3)			
6. & 10. Sowthistles combined			86.5	84.4	71.5	81.4
			(401)	(1696)	(5641)	(4832)
20. Winter oat	90.1	90.9	87.8	98	71.6	73.7
(*Triticum aestivum*)	(141)	(559)	(3736)	(9464)	(25249)	(20523)
21. Winter wheat	98	79.5	90.6	86.6	77.8	78.0
(*Avena sativa*)	(859)	(2581)	(8619)	(16159)	(31088)	(23460)
22. Oilseed rape	73.8	81.3	72.4	70.3	41.6	55.9
(*Brassica napus*)	(6603)	(12303)	(30799)	(77124)	(60763)	(47420)
23. Sugar beet	85.3	89.2	84.9	93.2	88.6	92.5
(*Beta vulgaris*)	(258)	(837)	(3149)	(8550)	(25439)	(23187)
	Average performance % (Weighted Average performance %)					
Plant Class Average	85.1	82.2	77.6	75.4	62.3	72.6
	(77.2)	(79.1)	(75.2)	(75)	(58.8)	(68.8)
Crop Average	86.8	85.2	83.9	87.1	69.9	75.0
	(77.1)	(81.8)	(77.9)	(76.8)	(63.2)	(71)
Weeds Average	78.2	76.2	74	70.8	59.3	71.7
	(78.2)	(64.4)	(63.3)	(68.4)	(51.2)	(65.1)

These outdoor results are lower than the ones obtained from the laboratory measurements, but are still at a level allowing substantial improvements in chemical reductions or efficient mechanical weed removal without affecting the yield. This performance reduction can be assumed to come from an increased number of uncontrolled variables in the classification. These variables can include leaf reflectance measurements at various angles, minor changes in illumination, or water or dirt on leaves. A greater number of classes also typically reduces the performance. All such factors can have a great effect on the final results.

The plant class average classification performance calculated directly from individual class performance figures decreased over time, while the weighted performance, as calculated with equation 6.2, remained relatively constant until 25 DAS. The weighted performance, $\bar{p}_w$, was calculated using

$$\bar{p}_w = \frac{\sum_i (p_i n_i)}{\sum_i n_i}, \tag{6.2}$$

where p_i is the performance figure of each individual plant type and n_i the corresponding number of windows classified in each case.

The stability of the weighted performance over time indicates that the classification method is robust in creation of classification functions, while the class average shows a better picture of how well individual classes can be classified and separated from each other. The weighted average results are greatly biased towards oilseed rape classification probabilities because the number of samples from this class was by far the highest in all but measurements gathered 25 DAS. A greater number of plant classes seemed to reduce the classification probabilities overall and generally crop plants were classified with a greater accuracy than weed plants.

It is difficult to find a reason why crop species classification is outperforming weed species classification when all species are classified separately. Weeds were classified better when a two class case was considered and presented in Table 14. It can be speculated that some biological differences in the plants make the difference. Crop plants have been developed over years while selecting the varieties that produce the best yield and most consistent plants. This is not necessarily true for the weed plants, although a weed plant is only a weed by definition at a given time in a given location. However, crop plant descriptors in the presented tests have on average 12% higher standard deviation than

weed plant descriptors.

Some selected scatter plots using the most descriptive wavelengths are shown in Figure 67 from 21 DAS training data. The separation of the data clouds in the figure represent how well different plant classes can be distinguished from each other. All weed plants are plotted in red and crop plants separately with different colours. Similar data representation was shown for laboratory measurements in Figure 48.
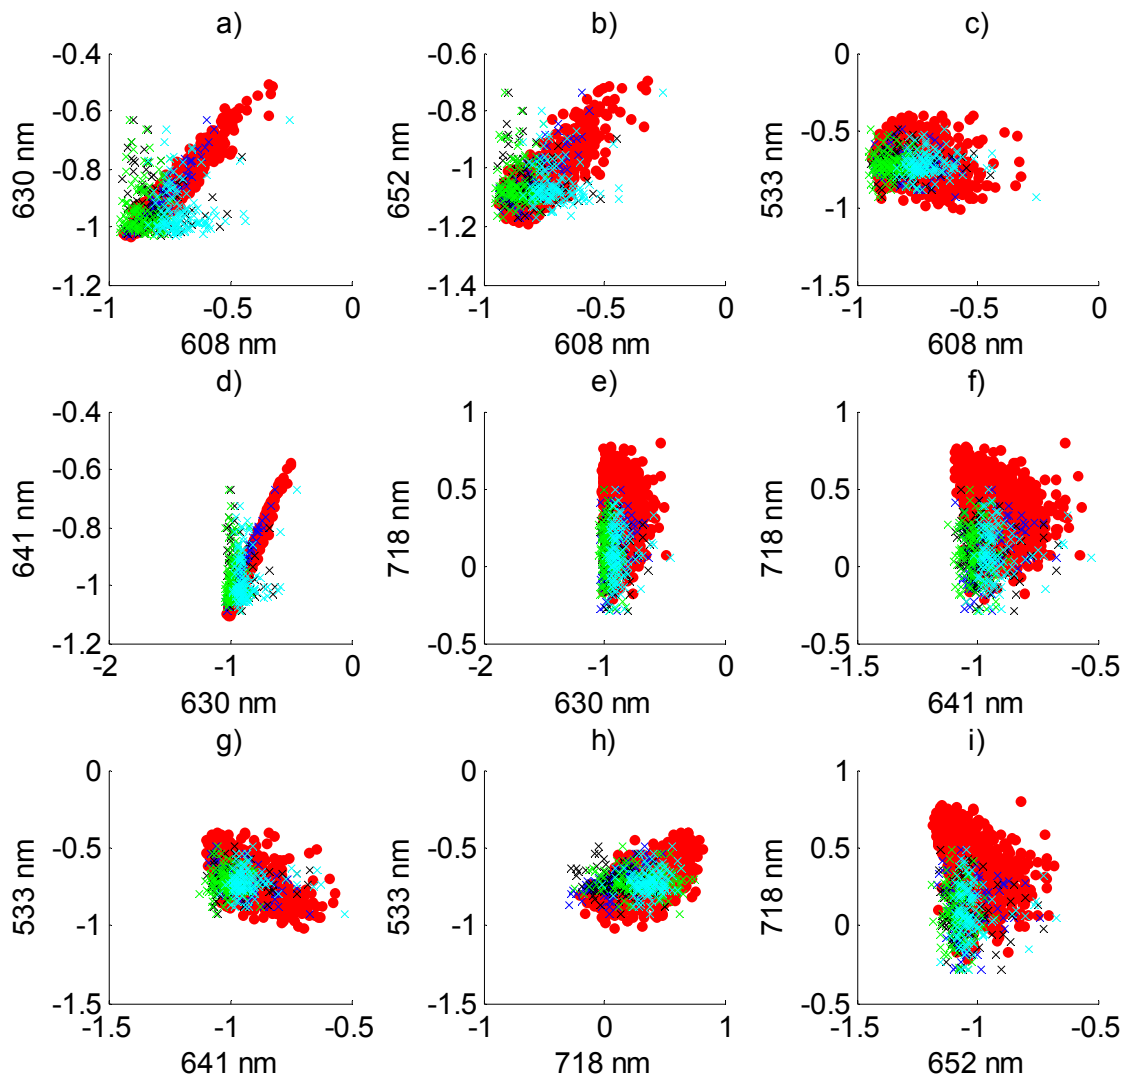
Figure 67 Scatter plots showing separation of crop and weed plants with selected descriptors from 21 DAS training data

The scatter plots illustrate the level of difficulty in classifying the plant species. Most data clouds are overlapping, but as the final classification results show, classes are separable using a large number of descriptors. Considerable amount of crop plant data points are separated from the weed plants in Figure 67 a), b) and d) and the crop plants are relatively distinct from each other. In general the combined weed plant data clouds are spread in

larger areas than the crop plant data clouds. The shown figures help in understanding why crop classification seems to work better than weed classification when all species are classified separately and why weed classification works better in a two class case of combined weeds and a crop plant.

The results in the Table 15 suggest that the ideal time for classification is any time from emergence until 21 DAS. This indicates that plants are easier to separate with the presented system when they are fairly small in size. This time constraint is well in tune with the ideal chemical application time shown in Figure 3.

Some visual examples of classification results and correctly and incorrectly classified windows at 21 and 25 DAS can be seen in Figure 68. a) - d) show examples from sugar beet bed and e) - h) from oat bed. Correctly classified windows are shown in grey and incorrectly classified windows in black in b), d), f) and h). The growth of the plants in four days can be seen when images a) and e) are compared to c) and g). At 25 DAS the soil was wet from rain and looked much darker than at 21 DAS. This did not effect the classification since plant segmentation was shown to work well in all conditions.

a) Sugar beet imaging area 11 21 DAS
b) ☐ correctly and ■ incorrectly classified windows

c) Sugar beet imaging area 11 25 DAS
d) ☐ correctly and ■ incorrectly classified windows

e) Oat imaging area 11 21 DAS
f) ☐ correctly and ■ incorrectly classified windows

g) Oat imaging area 11 25 DAS
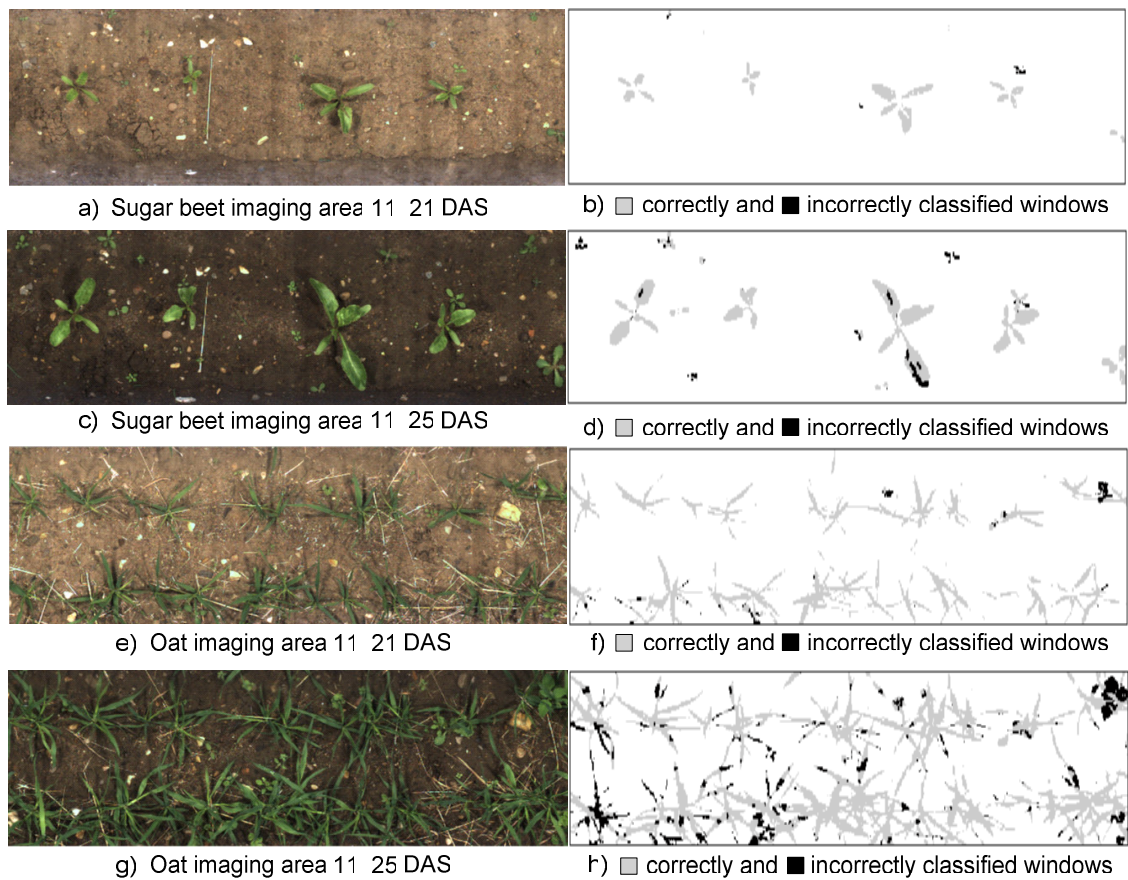h) ☐ correctly and ■ incorrectly classified windows

Figure 68 Examples of correctly and incorrectly classified windows within sugar beet and oat beds at 21 and 25 DAS.

A thorough investigation of the classification results and the number of false positives and negatives can be done with an error matrix, as shown during the laboratory test result analysis. An error matrix for the classification results gathered at 21 DAS is shown in Table 16. Here the distribution of the classification percentages for each plant class is detailed in each table column. The diagonal of the table shows the correctly classified cases in bold.

Table 16   Classification matrix for data gathered at 21 DAS. Percentage figures of correctly classified windows are shown in bold in the table diagonal. Class 12: the red dead nettle and 13: the common couch, other plant class numbers are as shown in Table 15.

Classified	Plant Class															
as (%)	1	2	3	4	5	6	7	8	9	10	11	20	21	22	23	6&10
1	**65**	0	0	1	3	0	2	5	0	0	0	0	0	1	1	0
2	1	**68**	5	7	0	0	4	1	2	0	0	0	4	0	0	0
3	0	4	**80**	1	0	0	0	1	6	0	0	0	1	0	0	0
4	1	3	2	**77**	0	1	4	2	3	2	0	0	0	0	0	1
5	4	0	0	1	**90**	1	0	2	0	0	0	0	0	0	0	1
6	0	1	0	0	1	**54**	0	0	1	31[1]	0	0	0	0	0	46
7	0	2	0	4	0	1	**67**	3	1	2	1	1	1	0	0	1
8	5	1	2	0	0	7	2	**50**	3	2	4	0	1	5[2]	0	5
9	1	1	4	2	1	1	0	4	**59**	0	4	0	3	3	0	1
10	0	0	0	0	0	29[1]	4	0	0	**58**	0	0	0	1	0	38
11	3	4	1	0	0	0	0	6	6	0	**68**	0	1	14[3]	5	0
12	1	2	3	3	1	0	2	3	11	0	3	0	1	1	0	0
13	1	0	0	3	0	0	0	8	1	3	0	0	0	2	0	1
20	7	0	0	0	0	3	0	4	0	0	0	**98**	0	0	0	2
21	2	5	0	0	3	3	9	1	1	1	2	1	**87**	2	0	2
22	5	4	1	0	1	0	0	10[2]	3	0	16[3]	0	0	**70**	0	0
23	5	5	1	0	0	0	4	0	2	0	2	0	0	0	**93**	0
6&10	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	**84**

[1] Misclassification high between two sowthistles (6&10). This is mainly due to mistakes in manual recognition. Column 6&10 represents the classification probability where the sowthistles are combined.

[2] Misclassification relatively high between crane's bill (8) and oilseed rape (22)

[3] Charlock (11) and oilseed rape (22) are very similar in early stages of their growth and are relatively difficult to tell apart. This is reflected in their classification probabilities.

From the Table 16 it is clear that sowthistle plant classes (6 and 10) are highly similar and its samples are being misclassified into each others' classes. This is mainly due to the mistakes in the manual recognition and creation of the ground truth data. A more representative result for sowthistle classification can be given by combining these two classes into one as is done in the rightmost column of the table. Misclassification between oilseed rape and charlock is also significant. In their early growth stages these plants are highly similar in appearance and also share near identical spectral signatures. This similarity and misclassification can explain the lower classification probability of oilseed rape

compared to the other crop plants. Exclusion of either oilseed rape or charlock from the tests would have increased the overall classification probabilities. This is evident when results from single plant beds are compared to the case with all possible plants. Exclusion of plants was thought counterproductive as it is important and meaningful to investigate and report the system performance using the most difficult cases available.

The results from the full data set can also be investigated as a two-class case of crop vs. weeds. This investigation is comparable to the results given with the individual plant bed analysis. The two-class results are shown in Table 17. The values were calculated from the full classification matrices, as shown in Table 16, including the correct and incorrect number of classified windows instead of classification percentages. All weed plant results and crop plant results were added separately together for final figures.

Table 17 Two-class classification performance over time, comparing weed classification against sugar beet and all crops.

DAS	Weeds	vs.	Sugar beet	Weeds	vs.	All Crops
11	99.3		89.4	83.4		84.1
14	92.5		91.0	68.7		99.5
18	95.7		87.4	74.2		80.3
21	97.6		93.7	83.1		78.5
25	95.9		90.8	83.8		71.6

The results in Table 17 show two cases; all weed species are compared with sugar beet or all crop plants. The results with sugar beet comparison are much higher than when all crop plants are included. This is again mainly due to the high misclassification probability between oilseed rape and charlock. The shown results with sugar beet are also higher than the ones obtained from the sugar beet bed in §6.3.1. This indicates that the classification method is not too sensitive to number of classes. The system seems much more sensitive to highly similar plants.

6.3.3 CLASSIFICATION OF INDIVIDUAL PLANTS

Individual plants were labelled in the ground truth images of oilseed rape and sugar beet. This information can be used to present classification results related to the individual plants. These results are shown in Table 18. A plant was considered to be correctly classified if the majority of the windows within the plant were correctly classified. The table also shows the total number of plants available for each class and the weighted average results for each day.

Overall, sugar beet plants (23) are classified with a great accuracy. Judging by the previous tables, oilseed rape (22) and charlock (11) classification are affecting each other.

Many of the individual weed plants (1-13) are not classified too well, but as shown in Table 17, the weeds are generally missclassified as other weeds rather than crop plants.

Table 18 Classification results for full plants. Correct classification is shown (%) with the number of available plants (#). The weighted average is calculated excluding classes 6 and 10 and including the combined class 6&10.

Plant Class	DAS									
	11		14		18		21		25	
	%	#	%	#	%	#	%	#	%	#
1					69.2	13	82.7	52	25.5	51
2							77.8	18	65.8	38
3					65.0	60	73.0	89	77.8	90
4							76.5	17	84.6	39
5							88.9	18	89.7	29
6					77.3	22	59.3	27	64.7	34
7							66.7	6	60.0	5
8			85.7	7	44.4	9	27.3	11	37.5	8
9					78.8	33	65.1	43	55.9	34
10					50.0	14	56.3	16	25.0	8
11	90.5	21	71.4	28	59.1	22	58.3	24	70.0	20
13					100	3				
6&10					88.9	36	90.7	43	73.8	42
22	76.1	209	82.2	202	85.5	186	85.8	176	50.8	65
23	92.5	40	92.6	54	98.2	55	100	53	98.0	51
Weighted Average	79.6		83.2		81.3		80.5		67.8	

6.3.4 EFFECT OF SPATIAL FILTERING

Spatial filtering of classified windows was shown to improve results in laboratory tests. The same filtering method was applied to the outdoor test data. All the presented classification figures are after the spatial filtering.

The effect of spatial filtering is visualized in Figure 69, where a) an original image and b) all classified windows are shown. The misclassified windows are shown in black in c) and d) before and after spatial filtering, respectively. As expected, the filtering step reduces the amount of individual misclassified windows and acts like a median filter reducing salt and pepper type noise in the images.

Figure 69 shows the filtering effect visually, but the effect was also quantified. The average improvement of the classification results from all classes and all data sets was 7.5%. Maximum improvement in any class was 21.6% and the minimum a decrease of 1.8%. An average improvement of 7.5% is significant and shows the effectiveness of the used spatial filtering technique.
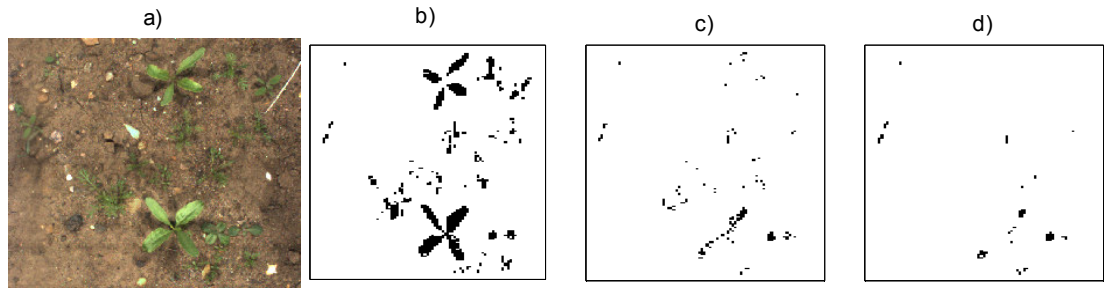
Figure 69 Illustration of the effects of spatial filtering in classified windows. a) Original colour image of sugar beet from location 43 taken 21 DAS, b) all found windows for classification, c) misclassified windows without and d) after filtering.

6.3.5 IMPORTANT DESCRIPTORS OVER TIME

Important descriptors for plant classification were evaluated during the laboratory testing and these results are given in §5.3.5. The laboratory tests did not include measurements over a period of different growth stages. The outdoor measurements were taken over a period of 2 weeks and therefore an additional analysis on the importance of the descriptors at different times can be done. These results will show if certain wavelengths and colour descriptors are consistently important for classification regardless of the growth stages of the plants.

Figure 70 shows the importance of the selected colour and spectral descriptors for each measurement day classification function separately. The combined results including all five measurement days are also shown. None of the single classification functions include all of the descriptors available. Relative descriptive powers for individual descriptors are taken from normalised F-values given directly by the stepwise discriminant process.
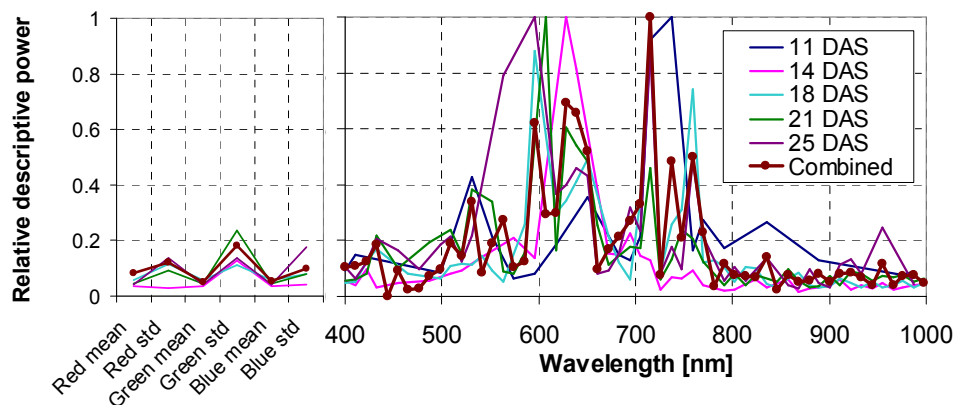
Figure 70 Important descriptors in classification at five different days

Similar wavelengths as seen previously in Figure 54 seem important for classification. The importance of these regions is relatively consistent throughout the data from different days. Wavelengths at 500 – 800 nm are the most important ones. Colour descriptors do not have high importance in the classification.

The general shape of the importance of the wavelength descriptors throughout the measurement range follows the CCD spectral sensitivity curve quite well. This can be considered coincidental since the results from the previous studies including other more sensitive measurement technologies back up the general outcomes of the analysis in this project.

The results strengthen the theory that certain wavelengths are consistently important for plant classification, and furthermore, their importance does not change at different growth stages. These results can be used to guide hardware selection and the system setup such that the properties of the system are optimal for the most important wavelength range.

6.3.6 TEMPORAL CHANGES IN PLANT SPECTRA

Dennison and Roberts (2003) have studied the changes in spectral canopy at different years and its effects on classification performance using satellite data. Unfortunately this and similar studies are not relevant for high resolution weed detection and the effects of temporal changes in spectral measurements over a period of some days. As explained in §2.2.5.2, the existence or creation of spectral vegetation libraries covering all possible variations in spectra is practically impossible. However, creation of libraries including local and short term spectral data could be achieved.

The measurements in this project were taken over a period of 11 to 25 DAS. These measurements can give an indication of the changes in plant spectra over a few days at early growth stages. The collected data represents spectral changes in local conditions where soil, nutrient and irrigation properties remained the same for all plants.

Figure 71 shows the mean spectra of oilseed rape on five different days over a 14 day growth period. Regions a), b) and c) are zoomed in for clarity. Some definite temporal spectral changes exist in this plant, although they appear random at most wavelengths. The order of mean values at different days over wavelengths changes. This variability can be explained by the local growing conditions and plant growth stages.

Investigation of the temporal changes at certain wavelengths in different plant types is shown in Figure 72. The plot for each plant a)-f) shows the trend of spectral change over five measurement days. The length of the vertical lines for each data point represents a value two times the standard deviation and the mean is represented by the horizontal lines connecting the data points on different days.
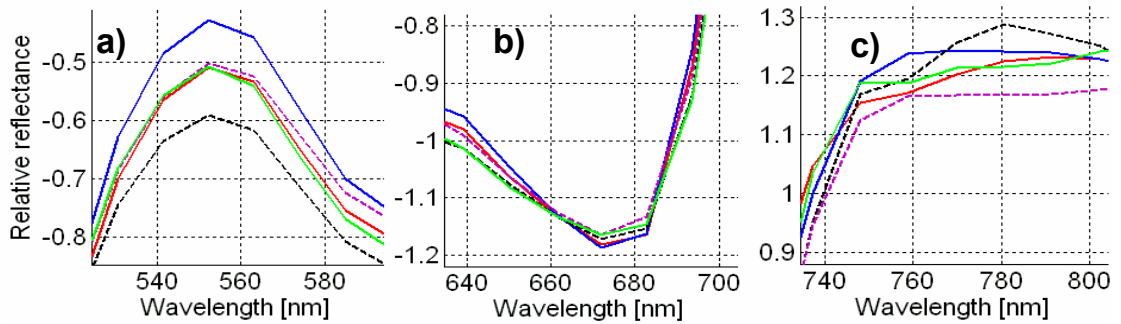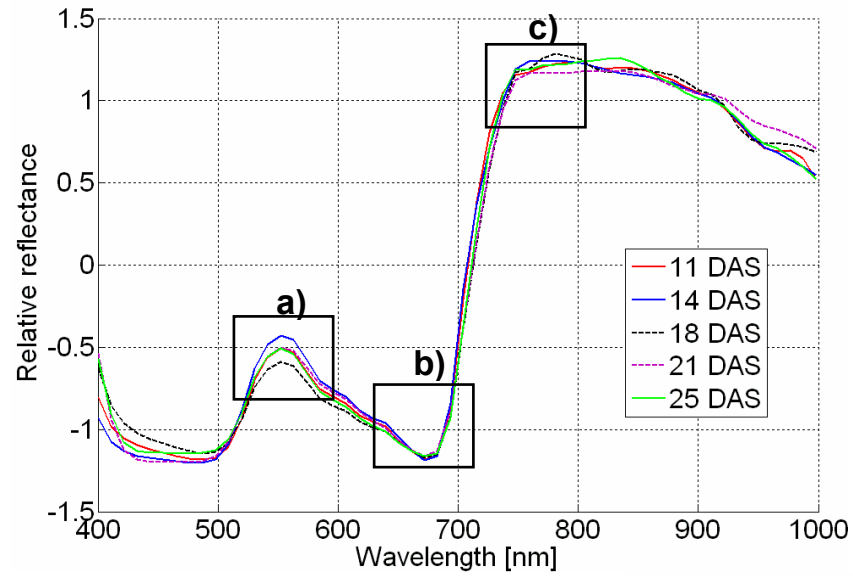
Figure 71 Oilseed rape mean reflectance at different days. a), b) and c) zoomed in regions of the total spectra
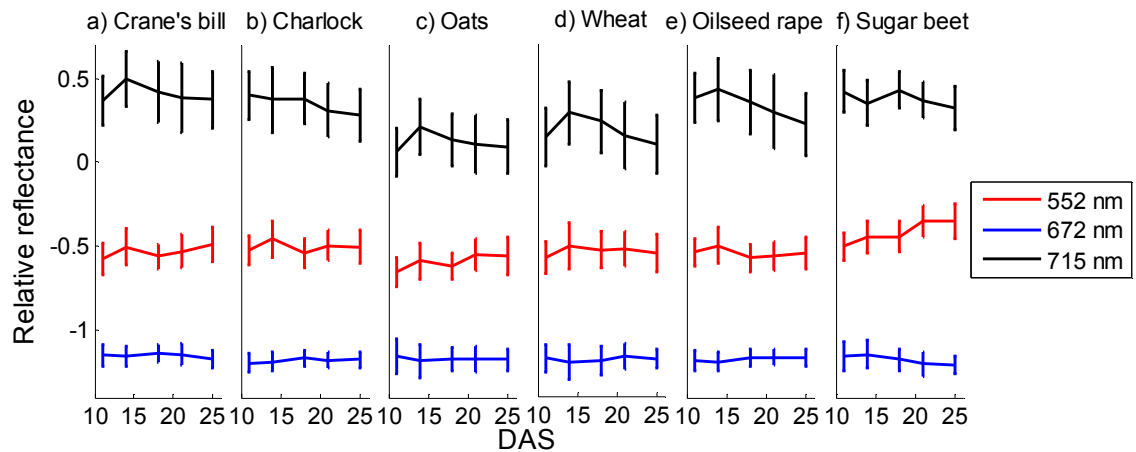
Figure 72 Temporal changes in three wavelengths for six plants

Clear spectral trends and reflectance levels over time for certain plants can be seen in this figure. For example, the grassleaved plants of oat and wheat share similar trends and reflectance levels at the shown wavelengths. This is also true for charlock and sugar beet data. This investigation can give an indication on how well the classification functions created from data from one day work to classify plants on different days. Unchanging spectra over several days indicates no potential problems for classifying a plant with

functions created from data gathered on different days and vice versa.

The information on temporal trend changes in spectra could potentially be used to create a model to predict spectral changes and thus make created classification functions adaptable and robust to temporal changes without gathering new local training sets. Estimation of such models with meaningful predictive powers would, however, require much more in depth knowledge of spectral changes due to varying parameters. Practical limitations in collecting measurements for all growing conditions and growth stages make the creation of thorough spectral libraries nearly impossible, although acceptable accuracies could be achieved by careful investigations.

6.3.7 ROBUSTNESS OF CLASSIFICATION WITH TRAINING SETS FROM DIFFERENT DAYS

It is interesting and important to examine how well the system can perform when the classification functions are created with data that is measured at different times. Examining the performance on all possible date combinations will determine if the system is capable of using training data including temporal changes. This is a type of sensitivity analysis of the system.

The classification performance was evaluated for training sets at each date against all other dates. Training sets from different days consisted of varying numbers of plant classes and naturally only the classes that are present in any particular training set could be used. The classification results of individual classes within any combination of training and sample sets were averaged together using the number of windows within each class as weights in averaging. The results were then normalized such that a value of one was given to the performance of a system with the same date for training and sample sets and other performance figures calculated relative to it. The results for all combinations of dates are shown in Table 19.

Table 19 Normalized average classification results for all combinations of training and sample sets. Diagonal values of 1 indicate the performance using training and sample sets from the same days. Other values show the normalised performance relative to this case.

		Sample set date (DAS)				
		11	**14**	**18**	**21**	**25**
	11	1	0.64	0.78	0.66	0.5
	14	0.37	1	0.9	0.54	0.54
Training set date (DAS)	**18**	0.06	0.55	1	0.36	0.35
	21	0.4	0.64	0.76	1	0.77
	25	0.21	0.37	0.52	0.53	1

It is clear that the performance is at its best when the data for the training set is gathered on the day of the sample measurements. The values to the right of the table diagonal on each row are the significant numbers indicating the predictive capabilities of the classification system created at an earlier stage. The figures to the left of the diagonal on each row show cases where the future properties of the plants would be known and used in classification. This could only be possible if a reliable spectral plant library would be available.

The results show that the system performance is sensitive to even short term changes. It is evident that the changes in reflectance properties discussed in §6.3.6 are significant for classification purposes during a course of just few days.

6.3.8 PERFORMANCE EVALUATION WITHOUT TRUE SAMPLES

It is valuable to be able to estimate the performance of the classification functions before any true samples are collected. This evaluation is then based solely on the information in the collected training set. The estimate would be important, for example, in a scenario where a reasonable training set has been collected from the field and an estimate in the performance would be required for a decision whether a larger training set is needed. Low performance at a particular plant class could indicate a need for a larger training set for that class.

The evaluation could be simply done by classifying the training set directly with the created functions. This self-estimation is considered positively biased and should never be used as a reliable result. This would be considered poor practice. It is unlikely that the underlying class boundaries and true variability could be reliably modelled using a limited training set.

Another approach, although only slightly more acceptable, is to use training set class mean and covariance matrices to create a simulated sample set. This is the method adopted in the discriminant analysis toolbox for Matlab by Vandev (2003). The methods provided in the toolbox create a simulated sample set with 6000 samples per class and classify this. A full classification matrix is given as an output. Table *20* shows the percentage difference between the simulated results and the actual results with real samples for all five data measurement days and available plant classes. Positive values in the table indicate that the performance with the simulated data was better and negative values indicate poorer performance with the simulated data.

Table 20 Difference between simulated classification results and the actual results. Positive values indicate the performance was better with simulated results. Values with a magnitude of 10 or over are shown in bold.

	Plant Class															
DAS	1	2	3	4	5	6	7	8	9	10	11	13	20	21	22	23
11											-2		6	-2	1	-4
14								**-11**			3		0	6	-8	**-11**
18	**33**		**10**		3			6	**-11**	3	**11**	-1	9	0	1	8
21	**14**	**10**	-2	-5	8	**13**	2	6	-2	7	-4		1	2	-6	0
25	**31**	7	5	**-11**	**17**	8	-9	-2	-6	**15**	-2		2	-3	**22**	-2

Generally the simulated performance matches well with the actual results. There are only a few cases with large differences in the performance values. The average over all the values is 3.1. Overall, the evaluation of the performance seems possible and meaningful with a simulated sample set as long as the limitations of the method are kept in mind.

6.3.9 CLASSIFICATION COMPARISON USING HIGHER SPECTRAL RESOLUTIONS

The optimal spectral resolution for the system was evaluated during the laboratory tests and explained in §5.3.2. The spectrograph internal focus was altered and the spectrograph field of view improved for data collection in the field tests. Such changes may have caused the system to perform better at different spectral resolutions than originally determined.

Most of the data in the field tests was collected with a spectral resolution of 10.7 nm and the only data set saved with a higher spectral accuracy was the one taken 25 DAS. Classification performance for this day decreased significantly from the previous days and was on average 62.3% with 10.7 nm spectral resolution data. Classification performance on this day was also investigated using data with spectral resolutions of 5.6 and 8.3 nm. Figure 73 shows the classification performances for individual classes using the three spectral resolutions. Class-wise values for 8.3 nm data can also be found in Table 15 in the column labelled "2[nd] 25 DAS".
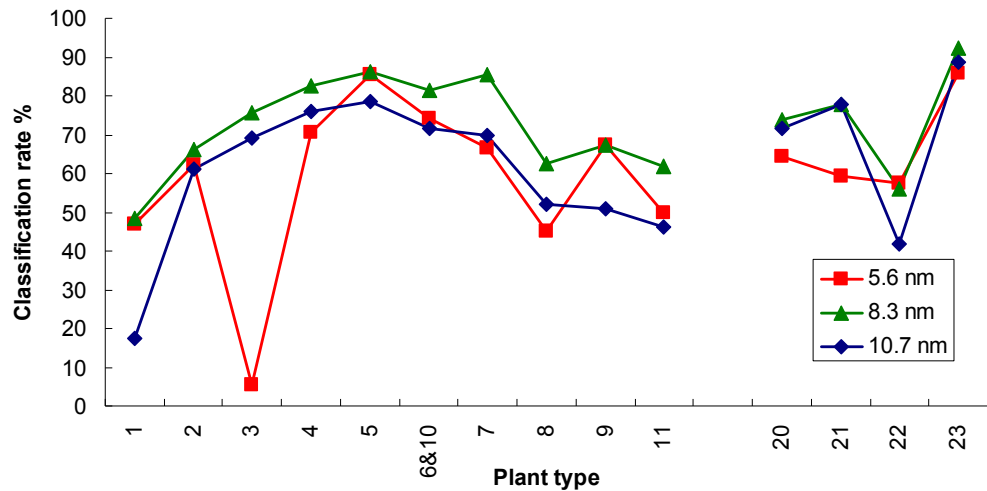
Figure 73 Classification comparison at different spectral resolutions 25 DAS

The average performance values over all plant types using 5.6 nm and 8.3 nm data were 60.1% and 72.6%, respectively. The highest resolution did not produce the best results. A significant improvement of over 10% compared to using 10.7 nm data was achieved by using 8.3 nm data. This implies that some improvements in classification accuracies may have been possible for the previous days as well. These potential improvements are impossible to quantify after the data was saved with a lower resolution.

The selection of spectral resolution seems sensitive to changes in imaging parameters. Therefore, it is suggested that an optimal resolution is evaluated for each different setup for best results.

6.3.10 PRINCIPAL COMPONENT ANALYSIS

The full measurement vectors for each window with 10.7 nm data consisted of six colour and 62 spectral variables. This is a large number of input variables for any classification system. Principal component analysis (PCA) can be used to create a set of uncorrelated variables, called principal components, from original data as explained in §2.2.5.1. The number of variables can be reduced by using only the most relevant principal components.

An analysis of the possible ways to reduce the number of variables was done for the measurement sets from all five days. The most important variables were first found by stepwise discriminant analysis and 22, 48, 55, 54 and 47 input variables were found for the five data sets from days 11, 14, 18, 21 and 25 DAS, respectively. PCA was then done for the training data sets using these variable sets.

An eigenvalue magnitude for each principal component represents the total variance explained by that principal component, i.e. if the sum of eigenvalues is 100 and the first eigenvalue magnitude is 20, the first principal component represents 20% of the total

variance in the data. Figure 74 shows a cumulative total variance proportion with different number of eigenvalues for all five data sets. The figure shows that a high number (>10) of eigenvalues and therefore principal components are needed to explain over 95% or even 90% of the total variance in the data sets. This indicates that the reductions in the numbers of variables will not be great if classification accuracy is to be maintained. When a threshold of 99% total variance is used, reductions of only 1, 7, 10, 10 and 7 variables can be made in the five data sets 11, 14, 18, 21 and 25 DAS, respectively.
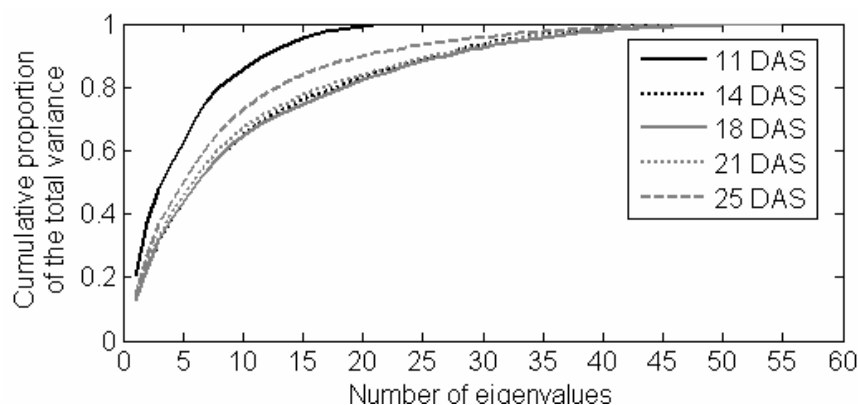
Figure 74 Cumulative proportion of the total variance vs. number of eigenvalues. Data plotted for all 5 measurement days.

Linear discriminant analysis was done using the reduced number of principal components for all five data sets. Table 21 shows the performance percentage difference between the classification results from the original discriminant analysis compared to the classification with principal components. Classification results for this analysis were taken before spatial filtering. The negative values in the table indicate that the PCA method had a lower performance than the original discriminant analysis and positive values indicate that the discriminant analysis performed better.

Table 21 Percentage difference in classification results between the PCA method and the original discriminant analysis. Negative values indicate that the PCA method had a lower performance. The values of under -10 are shown in bold.

DAS								Plant Class								
	1	2	3	4	5	6	7	8	9	10	11	13	20	21	22	23
11											-7		1	-2	-3	-6
14								-7			-4		-1	-1	**-11**	-9
18	2		-2			-5		-8	**-13**	-1	-3	0	-1	-7	-1	0
21	-5	-7	-6	0	-2	-3	-4	-3	-1	-3	-2		0	-3	0	-2
25	16	-3	0	-1	2	3	2	4	-2	-4	0		-1	-3	5	1

Overall the classification performance using the reduced number of principal components was lower than with the full set of variables selected with the stepwise

discriminant procedure. Since a great reduction in the number of variables could not be achieved and the classification performance was reduced when any number of variables was omitted, there was no reason to use principal components in classification with this particular system or data sets. In this case the benefit of classification accuracy from the full variable sets outweighs the benefits of improved calculation speed with the variable number reduction.

6.4 DISCUSSION

An explanation of field tests and a thorough analysis of the classification results using a dual camera system have been presented in this chapter. The system was shown to be capable of correct classification of all measured plants on average between 62.3 - 85.1% over different days in field conditions. In general, the classification results for crop plants were higher than for weeds. The average classification in a two-class case, where all the crop plants and weeds were combined together was up to 99.5% for crops and 83.8% for weeds. Similar results were shown when individual plant beds with one crop and some weed plants were analysed, and crop classification was up to 96.0% and weed classification up to 100% correct.

Borregaard, Nielsen et al. (2000) used a similar line scanning spectrometer as in this project and gained classification probabilities of up to 90% using only four plants. Vrindts, Baerdemaeker et al. (2002) also used the same spectral imaging technology with comparable number of total plants as used in this project. Two class sugar beet and weed classification produced results of up to 95% and 84%, respectively. Results for individual plant classes were not available and the spatial resolution was approximately eight times lower than in our system. Plant part locations could not therefore be given with similar accuracy as in this project.

Feyaerts and van Gool (2001) also show good results with a hyperspectral system classifying crop (sugar beet) and weeds with 80% and 91% correct classification, respectively. They give classification percentages for individual weeds but their final classification is for crop and weed only. The spatial resolution in their project was approximately ten times lower compared to the one used in this study.

The classification performance with the described system is comparable to these previously built systems. Classification performance figures collected from a set selection of plants can only give an indication of the true performance of the system, should the set of plants be different. As shown with oilseed rape and charlock, certain pairs of plants are harder to distinguish from each other than others. If charlock had been excluded from the

data sets the overall classification performance would have been much higher.

Short term temporal changes in plant reflectance spectra and its effect on classification had not been investigated previously. Measurements taken over a period of 14 days allowed this investigation. The analysis showed that the plant spectra changes significantly over short periods of time, and classification using training data from different days produces less accurate results. Modelling and prediction of these changes is difficult since the measurements do not include all possible variables in the growing conditions. The ideal solution is to gather a training set immediately before classification. The problem of gathering a training set autonomously or with minimal human input is not addressed in this project, but would have to be solved for usability of the detection and classification system.

Better classification results with the crop plants relative to the weeds may be due to the fact that the crop growth was more uniform than the growth of some weeds. The weeds had more variation in their size and growth stages. Therefore, a classification system trained with weed spectra that consists of plants at different growth stages and possibly also under different conditions is assumed to have a lower performance than a system trained with plants at the same growth stage and under same conditions.

The relative importance of wavelengths over 800 nm for classification compared to the band of 400 – 800 nm were low in these tests. This may be due to low CCD sensitivity or lack of powerful illumination beyond the visible spectrum. Since the importance of the higher wavelengths was lower, a lower cost monochrome camera sensor selection with a lower spectral response in the higher end of the spectrum could be made. A lower sensor resolution could also be allowed, since the spectral image was reduced before classification.

Principal component analysis with the measured data gave surprising results. It is common that data with high correlation between its variables can be reduced to only few principal components without losing much of the total variance. This was not the case with the data in this project. Significant reductions in the number of variables were not possible without losing the much needed plant class separability.

Overall, the results from the field test were promising and suggest that considerable reductions in chemical use or accurate mechanical weed control could be achieved with the proposed method. The presented results have shown the spatial and spectral accuracy limits needed for quality classification. These figures could be used to choose lower cost optical hardware and processing platforms aiming for real-time operation. Also, the system has potential to be further developed to be used with shape features, 3D plant analysis or in applications of detection and classification with other objects than living plants.

7.0 RECOMMENDATIONS & CONCLUSIONS

7.1 RECOMMENDATIONS FOR FUTURE WORK

Although automatic weed detection and classification has been intensively studied over the years, no commercial products have been introduced that can distinguish within-row weeds from any variety of crop plant. The main reasons for this are believed to be the economical limitations and the level of input required from a skilled operator. Operating speeds have been too low and hardware cost too high for the benefits gained when compared to current weed control methods. Skilled operator input has also been necessary to create a training set for the classification system. Creation of a real-time system capable of autonomous classification with minimal or no operator input has not been previously achieved.

Implementing the described system on a real-time platform is considered a task involving careful hardware selection and optimised algorithm code writing. FPGA boards could be used for fast processing both within the cameras and in an external processing unit. Relatively low-cost smart cameras with on-board processing capabilities and current interface technologies can reach high output rates of processed image frames estimated at over 100 fps.

Autonomous creation of the classification training set requires further research. This task is highly important in minimising the required human input. Feyaerts, Pollet et al. (1999) have investigated automatic clustering of crop and weed data using information on crop row structures. Their method was capable of correctly labelling all weeds and sugar beet at probabilities between 95% and 98.6% depending on the weed density. The within-row weeds were correctly labelled between 70% and 84.6%. No current system for automatic labelling of each crop and weed species separately exists. Also, further research is needed to adapt the weed and crop clustering to work with a variety of crop row structures and plant spacing.

An analysis on temporal changes in the reflectance spectra revealed that short term changes follow a trend when investigated at individual wavelengths. This information together with automatic clustering algorithms could allow creation of adaptable classification training sets. For example, when a classified measurement falls close enough to their respective class cluster mean in the training set, the measurement would be included in this training cluster with a time stamp. The oldest measurements in a cluster would then be removed, thus adapting the clusters to subtle temporal changes in the reflectance data.

Shape feature calculations were excluded from the window method based on heavy data processing needs and complications due to overlapping plants. Accurate leaf contour data could be used together with a recent shape classification method capable of rotation and scale invariant partial shape contour classification (Petrakis, Diplaros et al. 2002). Further research would be needed to extract contours from a single plant or leaf.

The presented imaging method creates a natural one-camera stereo imaging setup. If the movement of the colour camera between frames can be estimated accurately, a height estimate of the objects in view can be calculated. This height information could be a valuable addition to distinguishing objects. Calculating a height estimate for all pixels in an image requires heavy computing and its benefits for use with real-time weed detection applications should be carefully considered.

This section has outlined future tasks to make the presented system economically feasible and has given examples on how the system could be improved. The main problems to solve are the implementation of real-time operation and the minimisation of required human input.

7.2 CONCLUSIONS

The aim of this project was to create an optical plant detection system capable of robust classification of all plants in field conditions while providing accurate weed plant locations for mechanical or chemical weed control. For this purpose, a novel dual camera system consisting of a colour and a spectral camera was designed and tested in laboratory and outdoor field environments.

Spatially accurate imaging devices made a novel data processing method of window classification possible, where spectral and colour data were sampled and filtered in a grid-like formation while maintaining a spatial resolution of some millimetres. Classification of plants at this high spatial resolution has not been available before and locations of crop and weed plants could be estimated at a high accuracy of just few millimetres. The higher spatial accuracy of the presented system did not compromise the classification performance compared to previously built systems. Also, overlapping plants did not limit the use of the new system.

A weed detection system needs to be flexible enough to perform well with various plant combinations. This flexibility is one of the benefits of the presented two camera system. The classification figures of over 90% correctly classified windows in field conditions presented here are good enough for major reductions in the use of chemicals or for the precise mechanical removal of weeds. The presented system also has the potential

to include shape recognition and 3D height estimation of the scene into the classification by using data from a single colour camera.

The laboratory results using leaves from six tree and shrub plants with highly similar reflectance properties were comparable to the field tests when actual crop and weed plants were used. Although a direct comparison between the laboratory tests and the field tests is difficult due to differences in plant selection and external circumstances, the results indicated that a reliable system performance evaluation is possible with any type of plant.

The field measurements taken over a period of two weeks allowed analysis on short term temporal changes in the reflectance spectra of plants. It was shown that the changes in spectra over few days are significant for classification, and a training set created from measurements taken on a particular day cannot be reliably used after a few days. Also, their use on future crop cycles is thought limited due to subtle changes in growing conditions. The temporal changes follow a trend when investigated at single wavelengths, although the trends vary between the wavelengths. With further research, this information could be used to automatically adapt the old training set to the new reflectance properties. Investigation of classification performance levels on different days revealed that the optimal time for weed classification using the presented system is approximately 20 days after sowing.

The addition of a colour camera increases the system cost and data flow. In the presented system with window method data processing, the colour camera provides accurate information of leaf borders, is used for plant segmentation and RGB feature calculations. An estimate of the camera movement between the frames is also calculated using the colour data with means of cross-correlation calculations between image frames.

There is a certain amount of redundancy in the two camera system. For example, plant segmentation is possible at a slightly lower spatial resolution using the spectral data. The camera movement between frames cannot be estimated from the spectral data, but additional odometric or inertial sensors could provide the same information without the high processing loads from colour image cross-correlation calculations. The most important pieces of information for the window classification from the colour camera are the RGB features. In theory, the RGB mean data is also available in the spectral data. The use of a colour camera is well justified if shape, local texture, 3D information or accurate object borders are needed. The superior spatial accuracy and the possibilities this information provides are the benefits of the colour camera over the spectral camera.

The novel sensor system, window data processing and the temporal spectral analysis

are the most important original contributions of this project. The presented novel dual camera system was shown to be capable of correct plant classification up to an average level of 85.1% in outdoor field conditions with a selection of numerous weed and crop plants. Even performances of up to 100% were recorded when weeds were combined into one class. Intelligent window data processing methods were introduced to be used with the system. The analysis of temporal short term changes in plant reflectance spectra were shown to be significant for classification purposes. Real-time implementation of the system is thought possible with careful hardware selections and optimised algorithm code writing.

REFERENCES

Agronet (2005), Ohrasta oluen synty, Käsikirja mallasohran tuottajille, *Available at: http://agronet.fi/mallasohra/v6mallasohrankasvinsuojelu.htm*, Accessed: [25/1/2006].

Aitkenhead, M.J., Dalgetty, I.A., Mullins, C.E., McDonald, A.J.S. and Strachan, N.J.C. (2003), "Weed and crop discrimination using image analysis and artificial intelligence methods." *Computers and Electronics in Agriculture* 39(3), p. 157-171.

Alchanatis, V., Ridel, L., Hetzroni, A. and Yaroslavsky, L. (2005), "Weed detection in multi-spectral images of cotton fields." *Computers and Electronics in Agriculture* 47(3), p. 243-260.

Andersen, H.J. (2002). Outdoor Computer Vision and Weed Control. Computer Vision & Media Technology Laboratory. Aalborg, Denmark, Aalborg University, p. 152.

Antuniassi, U.R., Nery, M. da Silva, Carvalho, W.P.A., Ruiz, E.R.S. and León, M.J. De (2003), Performance Evaluation of an Optical Sensor for Weed Detection. ASAE Annual International Meeting, July 27-30, Las Vegas, Nevada, p. Paper number 031160.

Ashkenazi, V., Park, D. and Dumville, M. (2000), "Robot positioning and the global navigation satellite system." *Industrial Robot* 27(6), p. 419-426.

Åstrand, Björn and Baerveldt, Albert-Jan (2002), "An Agricultural Mobile Robot with Vision-Based Perception for Mechanical Weed Control." *Autonomous Robots* 13(1), p. 21-35.

Åstrand, Björn and Baerveldt, Albert-Jan (2004), Plant recognition and localization using context information. Mechatronics and Robotics 2004 – special session Autonomous Machines in Agriculture, September 13-15, Aachen, Germany, p. 1191-1196.

Åstrand, Björn and Baerveldt, Albert-Jan (2005), "A vision based row-following system for agricultural field machinery " *Mechatronics* 15(2), p. 251-269.

Azencott, Robert, Wang, Jia-Ping and Younes, Laurent (1997), "Texture Classification Using Windowed Fourier Filters." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19(2), p. 148-153.

Baerdemaeker, Josse De, Munack, Alex, Ramon, Herman and Speckmann, Hermann (2001). Mechatronic Systems, Communication, and Control in Precision Agriculture. IEEE Control Systems Magazine. 21(5), p. 48-70.

Bak, Thomas and Jakobsen, Hans (2004), "Agricultural Robotic Platform with Four Wheel Steering for Weed Detection." *Biosystems Engineering* 87(2), p. 125-136.

Billingsley, J. and Schoenfisch, M. (1995), "Vision-guidance of agricultural vehicles " *Autonomous Robots* 2(1), p. 65-76.

Blackmore, S. (2004), From Precision Farming to Phytotechnology. Mechatronics & Robotics '04, September 13-15, Aachen, Germany, p. 1182-1184.

Blaise, D. (2006), "Effect of tillage systems on weed control, yield and fibre quality of upland (Gossypium hirsutum L.) and Asiatic tree cotton (G. arboreum L.)." *Soil and Tillage Research* 91(1-2), p. 207-216.

Borregaard, T., Nielsen, H., Nørgaard, L. and Have, H. (2000), "Crop-weed Discrimination by Line Imaging Spectroscopy." *Journal of Agricultural Engineering Research* 75(4), p. 389-400.

Boyetchko, S.M. (1997), "Principles of Biological Weed Control with Microorganisms." *HortScience* 32(2), p. 201-205.

Brown, N., Jackson, M.R. and Parkin, R.M. (2000), "Automatic gravure print feature determination at production speeds." *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 217(8), p. 1101-1110.

Burks, T.F., Shearer, S.A., Heath, J.R. and Donohue, K.D. (2005), "Evaluation of Neural-network Classifiers for Weed Species Discrimination." *Biosystems Engineering* 91(3), p. 293-304.

Cho, S.I., Lee, D.S. and Jeong, J.Y. (2002), "Weed-plant Discrimination by Machine Vision and Artificial Neural Network." *Biosystems Engineering* 83(3), p. 275-280.

Clark, R.N., Swayze, G.A., Wise, R., Livo, E., Hoefen, T., Kokaly, R. and Sutley, S.J. (2007), USGS digital spectral library splib06a: U.S. Geological Survey, Digital Data Series 231, *Available at: http://speclab.cr.usgs.gov/spectral.lib06*, Accessed: [30/1/2008].

CRI (2008), VariSpec™ Liquid Crystal Tunable Filter, *Available at: http://www.cri-inc.com/products/components.asp?gclid=CJOEoKeilpECFQ81QgodlnygPg*, Accessed: [27/1/2008].

CTR (2007), Titech Visionsort, *Available at: http://www.ctr.at/carinthian_tech_research_english/news_presse/presseaussendungen/20070718_titech_en.php*, Accessed: [22/1/2008].

Dennison, P.E. and Roberts, D.A. (2003), "The effects of vegetation phenology on endmember selection and species mapping in southern California chaparral " *Remote Sensing of Environment* 87(2-3), p. 295-309.

Descour, M. and Dereniak, E. (1995), "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results." *Applied Optics* 34(22), p. 4817-4826.

Duda, R.O., Hart, P.E. and Stork, D.G. (2000), Pattern Classification Wiley-Interscience, ISBN: 0471056693.

Duval, J. (1997), Mechanical Weed Control in Cereals, *Available at: http://www.eap.mcgill.ca/Publications/EAP72.htm*, Accessed: [26/01/2006].

EdmundOptics (2005). Optics and Optical Instruments Catalog 2005, p. 232.

Feyaerts, F., Pollet, P., van Gool, L. and Wambacq, P. (1999), Hyperspectral image sensor for weed selective spraying. SPIE Conference on Advanced Photonic Sensors and Applications, Singapore, SPIE ISBN: 0277-786X/99.

Feyaerts, F. and van Gool, L. (2001), "Multi-spectral vision system for weed detection." *Pattern Recognition Letters* 22(6-7), p. 667-674.

Fyfe, S. (2004). Hyperspectral studies of New South Wales seagrasses with particular emphasis on the detection of light stress in Eelgrass Zostera capricorni. Wollongong, University of Wollongong, p. 389.

Gandhi, A. (2002), Content/Based Image Retrieval: Plant Species Identification, *Available at: http://web.engr.oregonstate.edu/~tgd/leaves/index.htm*, Accessed: [10/10/2006].

Gerhards, R. and Christensen, S. (2003), "Real-time weed detection, decision making and patch spraying in maize, sugarbeet, winter wheat and winter barley." *Weed Research* 43(6), p. 385-392.

Gerhards, R., Sökefeld, M., Timmermann, C., Kühbauch, W. and Williams, M.M. (2002), "Site-Specific Weed Control in Maize, Sugar Beet, Winter Wheat, and Winter Barley." *Precision Agriculture* 3(1), p. 25-35.

Gilchrist, J.R. and Hyvärinen, T. (2006). Hyperspectral Imaging Spectroscopy, A look at Real-Life Applications. The 2006 Photonics Handbook, p. 140-144.

Godwin, R.J. and Miller, P.C.H. (2003), "A Review of the Technologies for Mapping Within-field Variability " *Biosystems Engineering* 84(4), p. 393-407.

Granitto, P.M., Verdes, P.F. and Ceccatto, H.A. (2005), "Large-scale investigation of weed seed identification by machine vision." *Computers and Electronics in Agriculture* 47(1), p. 15-24.

Gregoire, T., Endres, G. and Zollinger, R. (1997), Identifying Leaf Stages in Small Grain, *Available at: http://www.ag.ndsu.edu/pubs/plantsci/weeds/w564w.htm*, Accessed: [10/12/2007].

Griepentrog, H. W., Nørremark, M., Nielsen, H. and Blackmore, B.S. (2005), "Seed Maping of Sugar Beet." *Precision Agriculture* 6(2), p. 157-165.

Hair, Joseph F., Black, William C., Babin, Barry J., Anderson, Rolph E. and Tatham, Ronald L. (2005), Multivariate data analysis, Pearson Prentice Hall, ISBN 0-13-032929-0.

Han, S., Zhang, Q., Ni, B. and Reid, J.F. (2004), "A guidance directrix approach to vision-based vehicle guidance systems." *Computers and Electronics in Agriculture* 43(3), p. 179-195.

Have, H., Nielsen, J., Blackmore, S. and Theilby, F. (2005), Autonomous Weeder for Christmas Trees - Basic Development and Tests, *Available at: http://www.mst.dk/udgiv/publications/2005/87-7614-868-8/pdf/87-7614-869-6.pdf,* Accessed: [1/3/2006].

Hemming, J. and Rath, T. (2001), "Computer-Vision-based Weed Identification under Field Conditions using Controlled Lighting." *Agricultural Engineering Research* 78(3), p. 233-243.

Hsu, S.M. and Burke, H.-h.K. (2003), "Multispectral Fusion with Hyperspectral Imaging Data: Detection and Classification." *Lincoln Laboratory Journal* 14(1), p. 145-159.

Huete, A.R., HuiQing, L. and Leeuwen, W.J.D. van (1997), The use of vegetation indices in forested regions: issues of linearity and saturation. IEEE International Conference on Geoscience and Remote Sensing, IGARSS '97. 'Remote Sensing - A Scientific Vision for Sustainable Development', 3-8 August, p. 1966-1968.

Hurle, K. (1997), "Concepts in weed control -- how does biocontrol fit in?" *Integrated Pest Management Reviews* 2(2), p. 87-89.

Johansson, M. and Baerveldt, A.-J. (2005), Segmentation of partly occluded plant leaves. Swedish Society for Automated Image Analysis (SSBA), Sweden, p. 77-80.

Johansson, M., Søgaard, H.T. and Baerveldt, A.-J. (2004), Robustness of Plant Recognition Based on Active Shape Models. Mechatronics and Robotics 2004 – special session Autonomous Machines in Agriculture, September 13-15, Aachen, Germany, p. 1185-1190.

Johnson, D.E. (1998), Applied multivariate methods for data analysts, ISBN: 0-534-23796-7.

Kerekes, J. (2006). Imaging spectrometers go commercial. LaserFocusWorld. 42(7), p. 63-68.

Keshava, N. (2003), "A Survey of Spectral Unmixing Algorithms." *Lincoln Laboratory Journal* 14(1), p. 55-78.

Kise, M., Zhang, Q. and Más, F. Rovira (2005), "A Stereovision-based Crop Row Detection Method for Tractor-automated Guidance." *Biosystems Engineering* 90(4), p. 357-367.

Koger, C.H., Bruce, L.M., Shaw, D.R. and Reddy, K.N. (2003), "Wavelet analysis of hyperspectral reflectance data for detecting pitted morningglory (Ipomoea lacunosa) in soybean (Glycine max)." *Remote Sensing of Environment* 86(1), p. 108-119.

Kristensen, E.F., Jørgensen, M.H. and Kristensen, J.K. (2005), Intra-Row Weed Control by use of Band Steaming. 15th IFOAM Organic World Congress, 20-23 September, Adelaide, Australia.

Kulcke, A. and Kerschhaggl, P. (2006), High Speed NIR Spectral Imaging in Recycling Applications. 3rd International Spectral Imaging Workshop, ECCV 2006, 13/5, Graz, Austria, p. 40-47 CTR Carinthian Tech Research AG.

Lawrence, K. C., Park, B., Windham, W. R. and Mao, C. (2003), "Calibration of a Pushbroom Hyperspectral Imaging System for Agricultural Inspection." *Transactions of the ASAE* 46(2), p. 513-521.

Lee, W. S., Slaughter, D. C. and Giles, D. K. (1999), "Robotic Weed Control System for Tomatoes." *Precision Agriculture* 1(1), p. 95-113.

Lee, W.S., Slaughter, D.C. and Giles, D.K. (1996), Development of a machine vision system for weed control using precision chemical application. Proceedings of ICAME-96, p. 802-811.

Leon, C.T., Shaw, D.R., Cox, M.S., Abshire, M.J., Ward, B., Wardlaw III, M.C. and Watson, C. (2003), "Utility of Remote Sensing in Predicting Crop and Soil Characteristics." *Precision Agriculture* 4(4), p. 359-384.

Levitt, D.G., Simpson, J.R. and Huete, A.R. (1990), "Estimates of surface soil water content using linear combinations of spectral wavebands." *Theoretical and Applied Climatology* 42(4), p. 245-252.

Lewis, M.M. (2001), Hyperspectral Discrimination of Vegetation - What is Possible. IEEE 2001 International Geoscience and Remote Sensing Symposium. IGARSS '01, December 3-5 2002, p. 2889-2901.

Llewellyn, R.S. and Powles, S.B. (2001), "High Levels of Herbicide Resistance in Rigid Ryegrass (Lolium rigidum) in the Wheat Belt of Western Australia." *Weed Technology* 15(2), p. 242-248.

Lobley, M., Reed, M., Butler, A., Courtney, P. and Warren, M. (2005). The Impact of Organic Farming on the Rural Economy in England, Centre for rural Research, p. 1-144.

LOT (2006), Lamp Spectral and Irradiance, *Available at: http://www.lot-oriel.com/site/site_down/ls_irradiance_deen.pdf*, Accessed: [24/1/2006].

Loux, M.M., Stachler, J.M., Johnson, W.G., Nice, G.R.W. and Bauman, T.T. (2007), Weed Control Guide for Ohio and Indiana, *Available at: http://www.btny.purdue.edu/pubs/WS/WS-16/WS-16.pdf*, Accessed: [22/2/2007/2007].

Manh, A.-G., Rabatel, G., Assemat, L. and Aldon, M.-J. (2001), "Weed Leaf Image Segmentation by Deformable Templates." *Agricultural Engineering Research* 80(2), p. 139-146.

Marchant, J.A., Hague, T. and Tillett, N.D. (1997), "Row-following accuracy of an autonomous vision-guided agricultural vehicle." *Computers and Electronics in Agriculture* 16(2), p. 165-175.

Martens, K. and Martens, M.-H. (2005), Blind cultivation, *Available at: http://www.newfarm.org/features/2005/0205/earlyweeds/index2.shtml*, Accessed: [26/1/2006].

Meyer, G.E., Neto, J.C., Jones, D.D. and Hindman, T.W. (2004), "Intensified fuzzy clusters for classifying plant, soil, and residue regions of interest from color images." *Computers and Electronics in Agriculture* 42(3), p. 161-180.

Micron (2008), PatchSpray, *Available at: http://www.micron.co.uk/product/patch_spray*, Accessed: [6/2/2008].

Mokhtarian, Farzin (2005), Search for similar shapes in the SQUID system: Shape Queries Using Image Databases *Available at: http://www.ee.surrey.ac.uk/Research/VSSP/imagedb/demo.html*, Accessed: [1/3/2007].

Mokhtarian, Farzin and Abbasi, Sadegh (2004), "Matching Shapes With Self-Intersections: Application to Leaf Classification." *IEEE Transactions on Image Processing* 13(5), p. 653-661.

Mokhtarian, Farzin and Mackworth, A.K. (1992), "A theory of multiscale, curvature-based shape representation for planar curves." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(8), p. 789-805.

Moran, M.S., Inoue, Y. and Barnes, E.M. (1997), "Opportunities and Limitations for Image-Based Remote Sensing in Precision Crop Management." *Remote Sensing of Environment* 61(3), p. 319-440.

Moshou, D., Ramon, H. and Baerdemaeker, J. De (2002), "A Weed Species Spectral Detector Based on Neural Networks." *Precision Agriculture* 3(3), p. 209-223.

Müller-Schärer, H., Scheepens, P.C. and Greaves, M.P. (2000), "Biological control of weeds in European crops: recent achievements and future work." *Weed Research* 40(1), p. 83-98.

NASA (2000), ASTER Spectral Libarary, *Available at: http://speclib.jpl.nasa.gov/*, Accessed: [30/1/2008].

NASA (2005), The Landsat Program, *Available at: http://www.earth.nasa.gov/history/landsat/landsat.html*, Accessed: [19/1/2006].

Neto, João Camargo, Meyer, George E. and Jones, David D. (2006), "Individual leaf extractions from young canopy images using Gustafson–Kessel clustering and a genetic algorithm." *Computers and Electronics in Agriculture* 51(1-2), p. 66-85.

Neto, João Camargo, Meyer, George E., Jones, David D. and Samal, Ashok K. (2006), "Plant species identification using Elliptic Fourier leaf shape analysis." *Computers and Electronics in Agriculture* 50(2), p. 121-134.

NTech Industries, Incorporated (2003), *Available at: http://www.ntechindustries.com/*, Accessed: [17/1/2006].

NTect Industries, Incorporated (2003), *Available at: http://www.ntechindustries.com/*, Accessed: [17/1/2006].

Ojala, T., Pietikäinen, M. and Harwood, D. (1996), "A comparative study of texture measures with classification based on feature distributions." *Pattern Recognition* 29(1), p. 51-59.

Ojala, T., Pietikäinen, M. and Mäenpää, T. (2002), "Multiresolution gray-scale and rotation invariant texture classification with Local Binary Patterns." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(7), p. 971-987.

Onyango, C.M. and Marchant, J.A. (2001), "Physics-based colour image segmentation for scenes containing vegetation and soil." *Image and Vision Computing* 19(8), p. 523-538.

Onyango, Christine, Marchant, John, Grundy, Andrea, Phelps, Kath and Reader, Richard (2005), "Image Processing Performance Assessment Using Crop Weed Competition Models." *Precision Agriculture* 6(2), p. 183-192.

Pedersen, S.M., Fountas, S., Have, H. and Blackmore, B.S. (2006), "Agricultural robots - system analysis and economic feasibility." *Precision Agriculture* 7(4), p. 295-308.

Pérez, A.J., López, F., Benlloch, J.F. and Christensen, S. (2000), "Colour and shape analysis techniques for weed detection in cereal fields." *Computers and Electronics in Agriculture* 25(3), p. 197-212.

Petrakis, Euripides G.M., Diplaros, Aristeidis and Milios, Evangelos (2002), "Matching and Retrieval of Distorted and Occluded Shapes Using Dynamic Programming." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(11), p. 1501-1516.

Phillipp, Isabelle and Rath, Thomas (2002), "Improving plant discrimination in image processing by use of different colour space transformations." *Computers and Electronics in Agriculture* 35(1), p. 1-15.

Prosilica (2006), High Performance Digital Cameras, *Available at: www.prosilica.com*, Accessed: [2/3/2006].

S.Arivazhagan, L.Ganesan and Kumar, T. G. Subash (2005), Texture classification using Ridgelet transform. Sixth International Conference on Computational Intelligence and Multimedia Applications (ICCIMA'05), August 16-18, p. 321-326.

SAS (2007), SAS, *Available at: http://www.sas.com/*, Accessed: [4/3/2007].

Scheepens, P.C., Müller-Schärer, H. and Kempenaar, C. (2001), "Opportunities for biological weed control in Europe." *BioControl* 46(2), p. 127-138.

Sclaroff, S. and Liu, L. (2001), "Deformable Shape Detection and Description via Model-Based Region Gropuing." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(5), p. 475-489.

Singh, A. (2006), Mechanical Weed Control in Organic Systems, *Available at: http://www.organicagcentre.ca/NewspaperArticles/na_mech_weed_ctrl.html*, Accessed: [26/1/2006].

Slaughter, D. C., Giles, D.K., Lamm, R.D. and Lee, W.S. (2000), Robotic Weed Control Systems for California Row Crops. EurAgEng, Warwick, UK, p. Paper No. 00-PA-018.

Slaughter, D. C., Lanini, W. T. and Giles, D. K. (2004), "Discriminating Weeds From Processing Tomato Plants Using Visible and Near-Infrared Spectroscopy." *Transactions of the ASAE* 47(6), p. 1907-1911.

Søgaard, H.T. and Heisel, T. (2002), Weed Classification by Active Shape Models. EurAgEng, Budapest, Hungary, p. Paper No. 02-AE-004.

Søgaard, H.T. and Lund, I. (2006), "Application Accuracy of a Machine Vision-Controlled Robotic Micro-dosing System." *Biosystems Engineering*(doi:10.1016/j.biosystemseng.2006.11.009).

Søgaard, H.T. and Olsen, H.J. (2003), "Determination of crop rows by image analysis without segmentation." *Computers and Electronics in Agriculture* 38(2), p. 141-158.

SoilAssociation (2005). Organic Market Report 2005 - executive summary, p. 3.

Sørensen, C.G. and Jørgensen, M.H. (2005), Intra-row Weed Control in Organic Crops – Technical Perspectives, Capability, and Operational Costs. ASAE Annual International Meeting, 17-20 July, Tampa, Florida, p. Paper Number: 054146.

Specim (1997), ImSpectror imaging spectrograph user manual Version 1.7,

Specim (2006), ImSpector, *Available at: www.specim.fi*, Accessed: [9/1/2007].

SPSS (2007), SPSS, *Available at: http://www.spss.com/*, Accessed: [4/3/2007].

Steven, M.D. (2004), "Correcting the Effects of Field of View and Varying Illumination in Spectral Measurements of Crops." *Precision Agriculture* 5(1), p. 55-72.

Steward, B.L. and Tian, L.F. (1998), "Real-time machine vision weed-sensing." *ASAE Paper No. 98-3033*.

Tang, L., Tian, L. and Steward, B. L. (2003), "Classification of Broadleaf and Grass Weeds Using Gabor Wavelets and and Artificial Neural Network." *Transactions of the ASAE* 46(4), p. 1247-1254.

Thenkabail, P.S., Enclona, E.A., Ashton, M.S. and van Der Meer, B. (2004), "Accuracy assessments of hyperspectral waveband performance for vegetation analysis applications." *Remote Sensing of Environment* 91(3-4), p. 354-376.

Thorp, K., Tian, L., Yao, H. and Tang, L. (2002), Development of Vegetation Indices for Hyperspectral Data. 2002 ASAE Annual International Meeting, July 28-31, Chicago, IL, p. Paper #021077.

Thorp, K.R. and Tian, L.F. (2004), "A Review on Remote Sensing of Weeds in Agriculture." *Precision Agriculture* 5(5), p. 477-508.

Tillett, N.D., Hague, T. and Miles, S.J. (2002), "Inter-row vision guidance for mechanical weed control in sugar beet." *Computers and Electronics in Agriculture* 33(3), p. 163-177.

Tominaga, S. and Okajima, R. (2000), Object Recognition by Multi-Spectral Imaging with a Liquid Crystal Filter. 15th International Conference on Pattern Recognition (ICPR'00), p. 1708-1711.

Vandev, D. (2003), Interactive Stepwise Discriminant Analysis in MATLAB. Pliska Studia Mathematica Bulgarica, Proceedings of the X International Summer Conference on Probability and Statistics and Seminar on Statistical Data Analysis, Sozopol, Bulgaria, p. 291-298.

VayTek (2007), Sony ICX285 Sensor Efficiency Curve, *Available at: http://www.vaytek.com/specroper.htm*, Accessed: [29/1/2007].

Vrindts, E. (2000). Automatic Weed Detection with Optical Techniques as a basis for Site-Specific Herbicide Application, Diss. de Agricultura Nr. 419. Faculty of Applied Bioscience and Engineering. Leuven, Belgium, Katholieke Universiteit Leuven, p. 146.

Vrindts, E., Baerdemaeker, J. De and Ramon, H. (2002), "Weed Detection Using Canopy Reflection." *Precision Agriculture* 3(1), p. 63-80.

Wright, G.G., Matthews, K.B., Cadell, W.M. and Milne, R. (2002), "Reducing the cost of multi-spectral remote sensing: combining near-infrared video imagery with colour aerial photography." *Computers and Electronics in Agriculture* 38(3), p. 175-198.

Xiang, H. and Tian, L. (2005), Automatic Camera Parameters Control Under Natural Lighting Conditions. 2005 ASAE Annual Meeting, 17-20 July, Tampa, Florida, p. Paper #053016.

Yao, H., Tian, L., Tang, L. and Thorp, K. (2002), Corn Canopy Reflectance Study With A Real-Time High-Density Spectral-Image Mapping System. 2002 ASAE Annual International Meeting / CIGR XVth World Congress, July 28-31, Chicago, Illinois, p. Paper #021076.

Zimdahl, R.L. (1993), Fundamentals of Weed Science, Academic Press, Inc., ISBN: 0-12-781060-9.

The mathematical functions in the Matlab discriminant analysis toolbox used in this project (Vandev 2003) needed to be calibrated to verify that the calculations were done correctly within it. This calibration of the linear stepwise discriminant analysis was done manually using small training and sample data sets shown in Table 22.

Table 22 Training and sample data sets for DA calibration

Class	Training Set				Sample Set			
	A	B	C	D	A	B	C	D
1	0.6823	1.8729	1.7210	1.6084	0.5192	2.6423	1.1120	0.4052
1	0.6488	1.5935	0.7961	1.5503	1.6356	1.9884	1.6699	0.4128
1	0.6598	1.8631	1.6923	1.1394	0.8193	2.0647	1.1255	1.5681
1	2.5127	2.0175	1.0132	0.8101	1.1479	2.6465	1.3001	1.1632
1	0.7217	1.9184	1.0435	1.1516	1.4056	2.6163	1.5200	1.1461
2	1.6897	1.7739	0.8438	2.0415	2.3607	1.4723	0.0359	1.4013
2	2.6197	1.8869	1.7616	0.7569	2.6509	2.1575	1.2532	1.7607
2	1.9786	1.9961	1.5786	1.0301	1.9107	1.3660	1.1148	0.8178
2	2.2217	1.2846	1.2290	1.3954	2.5455	0.8932	0.9052	0.7466
2	2.5383	1.4680	0.9729	1.1125	1.5513	2.5124	0.9718	0.6185
3	1.3829	2.0557	1.5324	2.9590	2.0393	2.5749	0.6091	1.5972
3	2.2125	2.1687	0.9151	2.1122	1.4260	2.2519	1.0989	2.6627
3	1.6630	1.8454	1.3535	1.9369	1.8198	2.0050	1.8887	2.0795
3	1.5051	1.2139	0.1057	2.4629	2.5006	1.6942	0.9461	2.0994
3	1.0053	1.8230	1.1228	2.6149	2.0424	2.2123	0.8966	1.5901

The data sets were created by randomly sampling from normal distributions illustrated in Figure 75. Class 1 seems separable with variable A, class 2 with variable B and class 3 with variable D. Variable C does not add value to the classification, but is used to show how poor descriptors are not included in the final classification functions.
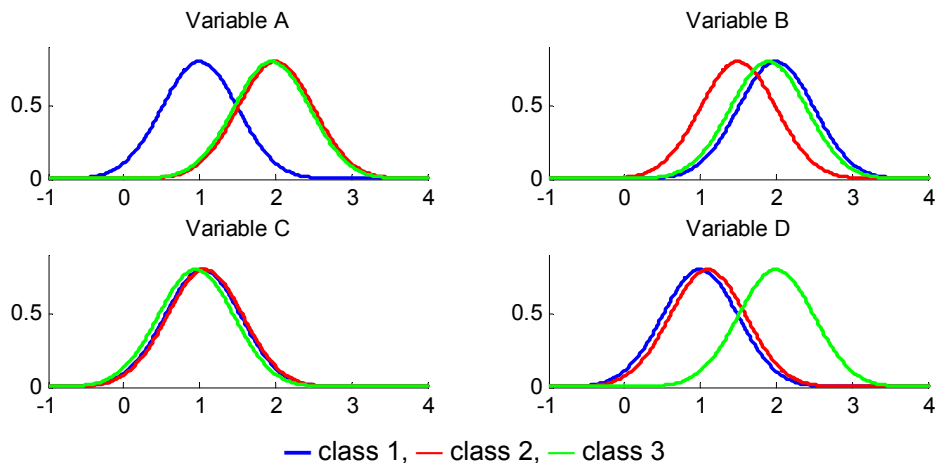

Figure 75 Normal distributions for the four variables used in calibration

Let us assume we have a training set consisting of vectors $(g_i, x_i), i = 1, 2, ..., n$, where g_i represents a group or a class a measurement x_i belongs to. Group means are then calculated as follows.

$$m(g) = \frac{1}{n(g)} \sum_i x_i \qquad (7.1)$$

Cross-products i.e. the within-group and between-group covariance matrices are given in the following two equations, respectively.

$$W(g) = \sum_i (x_i - m(g))(x_i - m(g))^T \qquad (7.2)$$

$$B_{in} = \sum_g W(g) \qquad (7.3)$$

The standard maximum likelihood-estimates are then calculated using

$$C(g) = \frac{1}{n(g) - 1} W(g), \qquad (7.4)$$

and

$$C = \frac{1}{n - G} B_{in}, \qquad (7.5)$$

where G is the number of groups or classes and n the number of samples.

In a linear discriminant case it can be assumed that within-group covariance is equal to between-group covariance and B_{in} can be replaced by $W(g)$ in equation 7.5. C also represents the total covariance of the system.

f-value, a measurement of difference between individual distributions, was calculated individually for variables entering the model and for variables to be removed from the model with the following equations:

$$f_{enter} = \frac{(n - a)(1 - W(g)/T)}{(W(g)/T)(b - 1)} \qquad (7.6)$$

$$f_{remove} = \frac{(n - a + 1)(W(g)/T - 1)}{(b - 1)}, \qquad (7.7)$$

where n is the number of training cases, a the number of total variables and b the number

of variables with a positive within-class covariance value. In this example, f-value of 2.0 was used for a variable to enter or to be removed from the discriminant model. Two sets of covariance matrices are being updated; one including the variables in the model and one including the variables out of the model. Updates of covariance matrices and variable *f*-values were required every time a variable was added or removed from the two models.

All of the shown equations were extracted from the toolbox functions and were manually used to calculate stepwise discrimininant analysis with the data in Table 22. The initial situation of the model with all variables out of the actual classification model with *f*-values is shown in Table 23. As expected from investigation of the data in Figure 75, variables D and A have the highest significance for separating classes from each other.

Table 23 Initial f-values for variables

Variable	*f*-value
D	7.9936
A	3.0677
C	0.3230
B	0.3049

In the first step, variable D was selected to enter the model. It had the highest f-value and it was over the set limit of 2.0. Within- and between-class covariance matrices and *f*-values were updated accordingly and the results for steps 1 and 2 are shown in Table 24. The results show how the variable A had the highest significance (> 2.0) after step one and was the next variable to be included into the classification model. After addition of the variable A in step 2, no variables had high enough *f*-values to enter the model or low enough *f*-values to be removed from the model. This was the end of the stepwise selection process and the final total covariance matrix of the model was then

$$C = \begin{bmatrix} 0.1708 & -0.1569 \\ -0.1569 & 0.3395 \end{bmatrix}.$$

Table 24 Two steps of the discriminant selection process showing variables in the classification model, model total covariance C and variables not in the model

Step 1	In model	f-value	Step 2	In model	f-value
	D	13.08		D	48.217
				A	22.778
	Not in model			Not in model	
	A	10.44		B	1.527
	B	0.7134		C	0.643
	C	0.0090			

The created model was then used to classify the sample cases using the knowledge of

class means *m(g)* and total covariance matrix *C*. Mahalanobis distance, as calculated using the following equation, was used to measure sample measurerent distances to group means using the variable correleation information availbale in the covariance matrix.

$$d(x,g) = (x - m(g))^T C^{-1} (x - m(g)) \tag{7.8}$$

The classification result was then determined by selecting the shortest distance to a class mean for each sample. Table 25 shows the classification results and individual Mahalanobis distances for the samples. All samples for class 1, 2 out of 5 for class 2 and 3 out of 5 for class 3 were correctly classified. Visual interpratation of the class overlaps in Figure 75 also indicate that samples in class 2 are the most difficult to classify correctly when only variables A and D are used.

Table 25 Classification results for the sample set showing individual Mahalanobis distances to group means

Sample	Mahal. distance to 1	Mahal. distance to 2	Mahal. distance to 3	Clasification result	True class
1	12.91	35.92	66.29	1	1
2	4.29	13.74	39.40	1	1
3	0.61	6.88	15.97	1	1
4	0.05	6.92	21.64	1	1
5	0.42	4.38	18.34	1	1
6	10.94	0.49	6.13	2	2
7	23.53	5.53	3.77	3	2
8	2.22	3.78	21.32	1	2
9	6.99	1.69	17.86	2	2
10	2.38	10.53	32.98	1	2
11	9.51	0.73	4.31	2	3
12	26.07	12.68	0.40	3	3
13	16.08	4.52	0.68	3	3
14	29.77	9.76	2.79	3	3
15	9.43	0.70	4.38	2	3

When the Matlab toolbox functions were used for calculations with the given data, the results were identical with manual calculations. The presented comparison of manual and tooldbox calculations gives the confidence to use the toolbox functions in further classification analysis.