University of Nottingham

Early detection of stress in strawberry plants using hyperspectral image analysis

by

Amy Lowe

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
School of Computer Science

October 2018

Declaration of Authorship

Declaration of Authorship I, Amy Lowe , declare that this thesis titled, 'Early detection of stress in strawberry plants using hyperspectral image analysis' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Abstract

Strawberry plants produce one of the highest quantities of soft fruits in the UK. The plants are grown in fields and glass houses where the environment is hard to completely control. There are a variety of biotic and abiotic stresses that affect the plants' production rate, lifespan and the quality of the fruit. It is possible to mitigate these stresses, but first they must be detected as early as possible.

Drought, or water deficit, is an environmental stress which impacts the plants' productivity. The stress can be monitored by looking for visible signs, but detection by a person can occur too late. Using technology could improve the detection before a person can see the signs. One such technology is hyperspectral imaging. Hyperspectral imaging has the potential to detect certain features in plants by examining their reflectance spectrum. In this work, a spectral range from visible to near-infrared will be used to record reflectance changes from plants during drought experiments.

The data collected in this thesis comprises strawberry plants undergoing drought conditions. Initial inspection of the data suggests that a difference in reflectance may exist throughout this time period, but the data is noisy and changes subtle: using whole-plant measures may not be as accurate as selecting individual leaves. Also, particular leaves may be more suited to hyperspectral inspection than others. Therefore, the work in the rest of the thesis focuses on automatically selecting suitable quality leaves from which to derive hyperspectral measures.

Leaf locations are selected automatically by combining 2D spatial information and select wavelengths from the hyperspectral data to extract leaf features (leaf centres and vein patterns), and this information is used to initialise a level set shape model to segment the leaves. The results from the 2D segmentation are evaluated, and results suggest clearly visible leaves can be segmented well automatically. However, it is hypothesised that further information is required to select the best possible leaves for analysis from the set of segmented leaves, so 3D plant information is then incorporated into the approach.

The orientation of leaves, and shadows, are likely to influence the spectral signatures, and occluded leaves are undesirable. Therefore, the leaf orientation, position and height are taken into consideration when selecting leaves for measurement. To do this, 3D surface models are created from multiple digital images captured at the same time as the hyperspectral data. The hyperspectral data is then mapped

onto the model and suitable leaves are identified based on their orientation, height and distance from the centre of the plant. This allows for selection of a subset of leaves as sources of hyperspectral measurement for each plant. This approach, from segmentation to leaf selection and measuring, is built into a fully-automated pipeline, and evaluated.

The resulting leaves are selected from the other time points in the time series data, and the spectral signatures are compared to the initial manual inspection. The automatically selected leaf data show similar spectral responses to the manual approach. The challenge of detecting changes in these spectral profiles is investigated and discussed.

Acknowledgements

Firstly I would like to start with my sincerest gratitude to my supervisors Dr Andrew French and Dr Nicola Harrison for all the guidance and support they have provided. Without their advice and continued support this thesis would not have been possible.

Andrew, thank you for always being there for discussions and guidance during these past four years and for your time and patience. I have really enjoyed this research project and your supervision of helping me when I needed it and for letting me explore different directions in the project.

Nicola, I really appreciate your help and support from the start of the PhD with the initial data collection, through to writing up. While I have predominantly been in Nottingham you have been willing to help. I really valued your guidance with the data collection methods and for your help over the past four years.

Thank you to everyone in the Computer Vision Laboratory for providing a friendly working environment. I have especially enjoyed my time sharing an office with Imanol, Shashank and Timur. They have been great friends, and always been there to have conversations with regardless of the subject, no matter how big or small.

I am also grateful to everyone at NIAB EMR who made me feel welcome in a new environment while I collected the data and for making me feel like I had never left whenever I visited.

I would like to thank AHDB, who partially funded this project. Without this funding the PhD would not have happened. The annual studentship conferences were especially helpful and interesting to see how research is used by commercial growers.

I would particularly like to thank my boyfriend Chris, for his love, encouragement and confidence in me.

Finally, I would like to thank my family who have always been there, especially my mum who always supports my decisions in life.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Strawberry plants (*Fragaria ananassa*) account for 1/7 of the soft fruit produced annually in the UK, 115 thousand tonnes were produced in 2015 and a similar amount in 2016 [1] [2]. Understanding how to manage and respond to stress is therefore important to a commercial strawberry grower.

Trying to predict stress in plants is complicated due to the amount of biotic and abiotic factors affecting the plants. Stress in plants can occur from a variety of sources. There is biotic stress which includes pathogens and parasites, and abiotic stress which includes natural changes in the physical environment (light intensity, drought). There are known visible signs to look for depending on the plant species and the disease/pest or drought under consideration.

A few examples of stress signs in the strawberry plants are described [3]. Powdery mildew caused by a fungus (*Podosphaera aphanis*) shows some or all of the following visible indicators: rolling of infected leaves, purplish or reddish blotches on leaves, the infected flowers and fruit will have a fine white covering of powder. Spider mites (*Tetranychus urticae*) are very small arthropod pests that infest the plants. The signs include bronze coloured leaves and eventually leaf death because the eggs are laid under the leaves and extract the sugar productions. If the infestation is severe it can reduce fruit production by up to 80% [4] A few visible signs of stress due to drought include curling and lighter coloured leaves (see Figure 1.1).

These stress signs are after some time has elapsed since the stress first occurred, before the visible signs appear there are changes happening in the leaf structure depending on the type of stress but this can not be visibly seen without using

technology to measure the leaf. Can early stress symptoms be detected through the use of a non invasive technology? This thesis will develop techniques to try to do that.

(a) Day 1 **(b)** Day 2 **(c)** Day 3 **(d)** Day 4

Figure 1.1: Images of a strawberry plant sensitive to water deficit under drought stress over time taken in the glasshouse at East Malling Research. By day 4 clear signs of stress are visible. Can the stress be detected earlier?

1.1 Motivation

Stress in crops can have an impact on the productivity of the crop and lead to yield loss. There are a few traditional ways of detecting stress caused by pests/disease (biotic) or environmental, eg. drought (abiotic stress) in strawberry plants. One way of detecting stress is manually checking the plants for visible signs of stress. This involves growers and agronomists regularly checking their plants for symptoms of disease or adverse conditions such as drought or nutrient deficiencies. This method is time-consuming and costly.

As an alternative, the health of a crop could be automatically measured to increase efficiency and reduce costs to the grower; however, at this stage easily visible signs of stress are measured rather than pre-symptomatic signs. Automated phenotyping is a high throughput screening of the plants, measuring the observable properties such as the leaf area and leaf pigment content. A plant's phenotype can be described as observable complex traits influenced by the genotype and the environmental conditions [5]. Automated phenotyping can be expensive due to the advanced technology and technical equipment needed, and subsequent analysis, although research is aimed at reducing the cost by using low cost cameras and phenotyping systems [6] [7]. Hyperspectral imaging is a recent addition to phenotyping technology, which can measure extra colour information that cannot be observed by eye or a standard digital camera.

Techniques for measuring the health of a crop also include vegetation indices. There are indices for detecting specific characteristics which combine the reflectance properties of two or more wavelengths; for example Normalised difference vegetation index (NDVI) is used to detect stress caused by a cereal pest in wheat [8]. There are a lot of vegetation indices but only a small number have been systematically tested [9]. This thesis will look more at hyperspectral imaging and analysis (including using indices) in the next Chapter.

Finding ways to automatically detect stress before it is easily visible would be beneficial to reduce potential losses to the quantity and quality of the strawberry crop. Combining image based phenotyping, hyperspectral imaging, and novel analysis methods may allow for the detection of early stress signs without manual labour cost. This is the motivation behind the work in this thesis.

1.2 Objectives

The overall objective of the thesis is to develop methods to detect the onset of stress in strawberry plants. The first question is: Can the stress be detected in Strawberry plants? This leads onto the next question: If so, how early can the stress be detected before visible signs appear? If the stress is detected early then the plant can be treated so the impact of the stress would be minimal to the plants productivity. Also it could potentially stop the spread of the cause to the whole crop. Before it can be detected, though, the thesis will examine the data processing and analysis required to take a good quality hyperspectral measure.

The main objectives of the thesis are as follows:

1. Dataset capture - capture novel hyperspectral datasets of plants undergoing stress.
2. Manual analysis of these datasets to try to identify stress onset indicatives.
3. Develop segmentation for individual leaf-based hyperspectral measures
4. Develop geometry compensation using 3D information to improve hyperspectral measurements
5. Create a pipeline to enable automatic hyperspectral measures for use in stress detection
6. Evaluate the pipeline with a view to early stress detection

In order to analyse plants in images they first need to be segmented and suitable measuring locations identified before the analysis can be applied. This could involve finding the whole plant, finding the individual leaves or finding the angles of the leaves, for example.

Image segmentation is widely researched and can be challenging, when it comes to plants it can be very challenging. Due to occlusions with overlapping leaves and similar colours, sometimes it is hard to define the boundary even by the human eye. Also plants vary in appearance; size of the plants, size of the leaves, colour, shape and complexity. Additionally shadows, leaf angles and lighting can also affect the appearance of the plants in the images, making plant image analysis a particular challenge.

Hyperspectral data can be particularly sensitive to such effects; by imaging plants in close proximity to the camera this introduces 3D effects for the hyperspectral data and a new challenge. The angles of the leaves and shadows can affect the data [10]. How much does this affect the detection of drought? Also how do you select the desired leaves on the plant?

There are different stress symptoms in plants depending on the biotic or abiotic stress involved and the aim is to determine whether it is possible to detect indicators of stress before these signs are visible to a human. Hyperspectral imaging will be used to analyse the spectrum range from the visible to near infrared wavelengths during the onset of the particular stress, and analysis methods will be developed to help gain insight from the hyperspectral data.

1.3 Problem domain: Hyperspectral Imaging for plant phenotyping

Hyperspectral imaging for crop health started with remote sensing, however in recent years the technology has become affordable to bring the research into the laboratory and glasshouse. Here the technology and data underpinning hyperspectral imaging is briefly introduced.

1.3.1 Imaging technology

There are various hardware approaches behind hyperspectral imaging spectrometers, which means there are different ways that the image is captured. In one example, the incoming light passes through a convex grating (or a prism) which separates the light into narrow wavelengths. This separation is then recorded on a light sensitive chip (similar to a standard digital camera). Examples of operation include push broom, filter wheel, liquid crystal tunable filters amongst others [11]. A pushbroom device, taken as a typical example, has three components; the camera, a spectrometer and a lens. This system simultaneously captures a single spatial line of the image, and the whole colour spectrum range. Then the camera or object is moved and the next pixel line is captured (the broom is 'pushed' forwards, hence the name), effectively making the camera a line scanner, with the final image being built up after the full scan is complete. An alternative to pushbroom is a snapshot approach, where the entire image is captured at once. To date, pushbroom technology has seen the most use, but recent advances in snapshot technology are increasing the uptake and possibilities related to phenotyping and analysis. In this project, pushbroom hyperspectral cameras are used.

1.3.2 Hyperspectral data

This thesis will require a good understanding of the nature of hyperspectral data. A hyperspectral image contains both spatial and spectral information, which can be represented in three dimensions. The spatial information is the pixels location along the x- and y-axis, and the spectrum is represented along the z-axis. This can be visually displayed as a *hypercube* (Figure 1.2b) where white represents high reflectance and black is low reflectance. Each band of the spectrum can be displayed as a greyscale image which displays the intensity of the light. Figure 1.3a shows a select number of greyscale images of the same plant in the spectral range of 400nm to 1000nm.

(a)

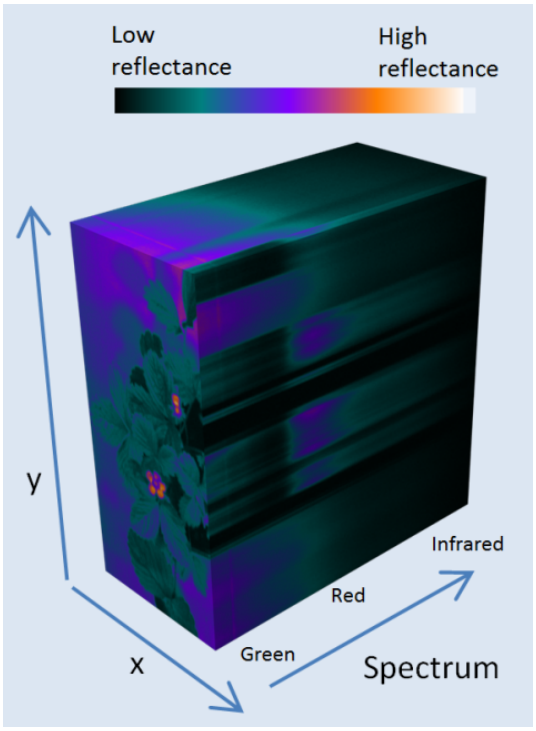

(b)

Figure 1.2: (a) is the RGB image of the Hawaii strawberry plant where the red box indicates the section that has been imaged. (b) is the hypercube of the selected area. The x and y show the pixels in the spatial location and the spectrum is respresented in the z axis from green to near infra-red wavelengths

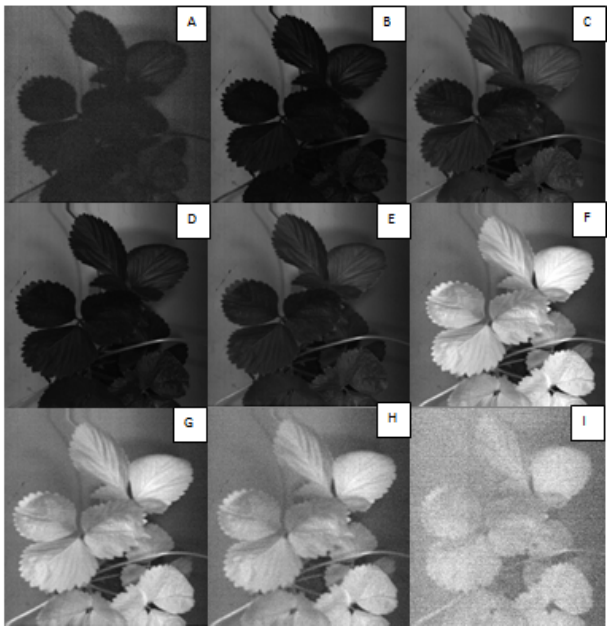

(a)

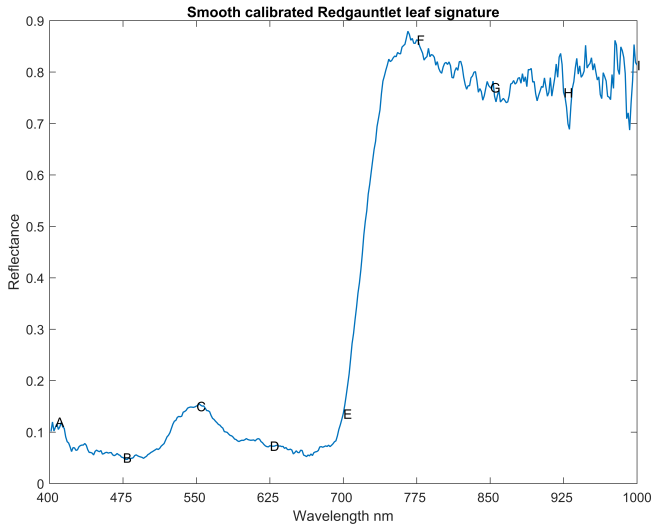

(b)

Figure 1.3: (a) Select greyscale images over the spectrum of a strawberry plant. The hyperspectral image is displayed as greyscale images at selected band intervals over the spectrum where the greyscale images are the intensity of the light at a specific band: A is 400nm, B is 472nm, C is 546nm, D is 620nm, E is 694nm, F is 768nm, G is 842nm, H is 916nm and I is 990nm. (b) is the spectral signature of the plant where the letters A-I correspond to the images in (a)

Healthy leaves and vegetation look green to us because the leaf pigments absorb most of the blue and red wavelengths ($\sim$ 450nm and $\sim$ 660nm) while most of the green wavelengths ($\sim$ 550nm) are reflected back. The red and blue wavelengths

are absorbed by chlorophyll. Chlorophyll absorbance has two peaks at $\sim$ 450 and $\sim$ 660nm. When the leaves become unhealthy there is less chlorophyll content, therefore some red wavelengths are also reflected instead of being absorbed, which is why the leaves start to look yellow. In the infrared range the wavelengths are longer which allows them to travel through the leaf to the cell structure, specifically the spongy mesophyll [12]. More light is reflected back in the infra-red range than in the visible range of the spectrum for vegetation; however the human eye cannot see this. Figure 1.4 shows the spectral signature of a healthy strawberry plant. There is high reflectance in the infrared points (Figure 1.3 G,H,I).

There are some common features of hyperspectral plant data that can be explained with reference to Figure 1.3b:

Visible light (A-D) are selected points in the visible light of the spectrum. C is the small reflective peak at the green wavelength of the spectrum, representing the green colour we perceive leaves as.

Red edge (E) is the wavelength at the steep reflectance incline called the 'red edge', this is a narrow section in the electromagnetic spectrum (690-740nm) where the visible spectrum ends and the near infrared starts. This section has a large change in spectral response space for green plant material, since chlorophyll strongly absorbs wavelengths up to around 700 nm, and hence the material has low reflectance in this range, but it is strongly reflecting the infrared (from about 720 nm).

Near infra-red (F-I) is the near infra-red section of the spectrum, typically 0.75 - 1.4um. In infra-red very little energy is absorbed by the leaves and therefore most of the energy is reflected.

Spectral signature A-I is the complete hyperspecral signature of a plant in the visible and near infra-red sections of the spectrum. This data looks continuous but they are actually individual contiguous wavelength band samples (0.74nm each) along a spectrum; in total 812 wavelengths are displayed in the graph. This is what the data looks like when it is captured compared to for example a vegetation index which may just use two wavelengths to determine the health of the plant or crop, selected from this continuous profile.VIs will be considered in more detail later in the thesis.

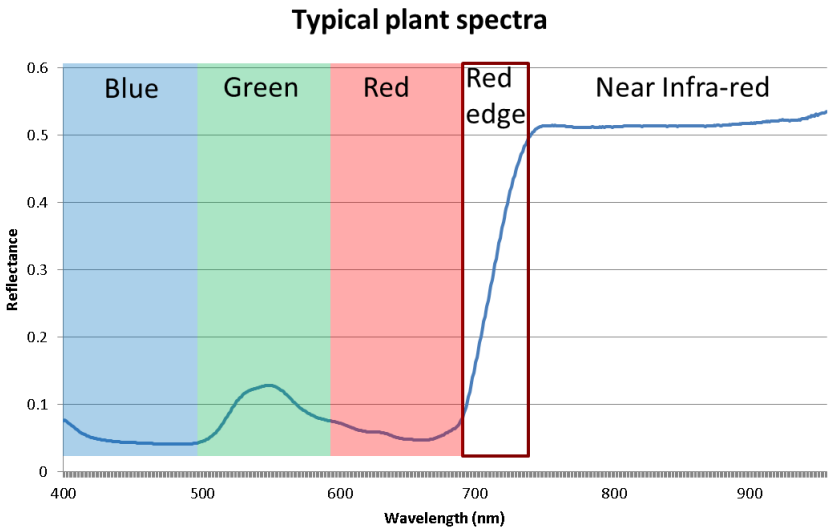

Figure 1.4: The spectral signature of a strawberry plant's leaves. The blue, green and red sections refer to the perceived colouring in the visible range, and the red box represents the red edge range. The plot is the reflectance of the light from the plant over the spectrum and this is known as a spectral signature. This data covers the range from 400nm to 1000nm. Each band width is 0.74nm.

More details on hyperspectral imaging and data will be presented in Chapter 3.

1.4 Thesis organisation

Chapter 2 - Literature review This Chapter is a general literature review, the first section outlines computer vision techniques for segmenting plants. The second section features an overview of different application to detect stress in plants [13].

Chapter 3 - Data collection and initial manual analysis explains the acquisition and processing of the datasets including calibration and identifies the challenges associated with the datasets. Also in this Chapter the data is preliminarily analysed using a manual software tool to select data measures and compare the spectral signatures. This is followed by experimental application of a clustering technique to separate the materials of the image. (Objectives 1 and 2)

Chapter 4 - A model based approach for segmenting leaves in hyperspectral images Development in this Chapter leads to segmenting individual leaves in 2D hyperspectral data with the aim to begin selecting certain leaves to be analysed. This process includes a new leaf model pipeline for segmenting the

leaves on a plant (Objective 3).

Chapter 5 - 3D plant reconstruction and data mapping involves finding extra information from the images using depth and/or 3D to help overcome the challenges of hyperspectral data. There are two different approaches, first depth is estimated for the images using stereo pairs of hyperspectral data. Then several digital SLR images are used for 3D reconstruction and surface reconstruction of the plants via multi-view reconstruction. The hyperspectral data is then registered onto the 3D model in order to ascertain the leaf properties for each hyperspectral data point (Objective 4).

Chapter 6 - Using 3D information to select leaves from which to measure hyperspectral data This Chapter creates an automated process to identify suitable leaves from which to take measures.

Chapter 7 - Hyperspectral data analysis of the selected leaves over time The Chapter analyses the time series data that has been selected through Chapters 4-6 for the early detection of stress in strawberry plants (Objective 5). The time series for the automatically selected leaves (Chapter 6) and manually selected leaves (Chapter 3) are compared.

Finally **Chapter 8 - Conclusions** concludes the thesis with a summary of the research presented and gives an overview of the possible directions for future work.

1.5 Contributions

The main novel contributions of this thesis are:

1. A review of the emerging field of using hyperspectral imaging to detect stress.(Chapter 2)
2. A leaf segmentation approach using shape models seeded from local normalisation and leaf vein detection. (Chapter 4)
3. A pipeline for mapping 2D hyperspectral data onto a recovered 3D plant model. (Chapter 5)
4. A novel data analysis approach to try to detect drought stress given 3D-mapped hyperspectral data, selecting automatically the best leaves from which to take measurements. (Chapter 6)

5. A comparison of this new method with manual stress detection methods. (Chapter 7)

The second contribution involves merging existing techniques together to locate and segment leaves in images and is covered in Chapter 4. Plant images are complex with occlusions, shape, colour intensity and shadows contributing to change in the appearance of the leaves. Local normalisation is usually used to normalise the intensity of images however this technique has been adapted to locate leaves. The Hough transform is then used to find the veins. A leaf template is translated and rotated according to the seed centres and vein direction. Then level sets with a shape prior (the placed shape template) is used to actually segment the leaves.

The third contribution is explained in Chapter 5 and involves registering the hyperspectral data onto a 3D model. This process is a pipeline which includes 3D reconstruction and surface creation before the registration process. This is a different approach to collecting full 3D hyperspectral data directly, that requires less data and is potentially a viable option to automatically collect the data in a glasshouse in the future. This process just needs *one* hyperspectral image and typically 3-7 RGB images of the plant from different angles. With the limited amount of images and viewing angle required this could be replicated with a simple capture process on a rig in the glasshouse or field in the future.

The fourth and fifth contributions are described in Chapters 6 and 7, this involves selecting data from the 3D models for further analysis. The data is selected automatically using several features from the plant: The leaf angles, the leaf height, the leaf distance and, also potentially the flatness of the leaf. These features are combined and weighted to select the best three leaves per plant. The data selected is then compared in Chapter 7 with the manually selected data in Chapter 3.

Additional contributions are the hyperspectral imaging protocols and datasets for plants experiencing 3 stress conditions (Chapter 3).

1.6 Publications

Various Sections of the thesis have been presented in various forms:

- 'Hyperspectral image analysis techniques for the detection and classification of the early onset of plant disease and stress' - Plant methods 13, no. 1

(2017):80 [13] - Hyperspectral literature review (Chapter 2.2)

- 'Meaningful region segmentation using hyperspectral imaging' - Image Analysis Methods for the Plant Sciences Workshop, Angers, France, 2016 - Presentation/Poster (Chapter 4)
- 'HyperPatches - a pipeline for reconstructing 3D plant models with registered hyperspectral data' - ICCV conference workshop on CVPPP, Venice, Italy, 2017 - Poster - (Chapter 5)
- 'Hyperspectral imaging with strawberry plants - Mapping hyperspectral data onto a 3D plant model' - Image Analysis Methods for the Plant Sciences Workshop, Nottingham, UK, 2018 - Presentation/Poster (Chapter 5)
- On the right wavelength - Article in The grower, page 25, issue Oct/Nov 2017 [14]

Chapter 2

Literature review

This Chapter reviews related work and is divided into two Sections; (2.1) Computer vision and image segmentation methods applied to plant data, (2.2) Techniques for detecting biotic and abiotic stress in plants.

The first Section is related to prior segmentation work used on similar colour and hyperspectral datasets that are viable options for finding the plants and leaves in the image data. The second Section describes and discusses prior approaches that have detected biotic and abiotic stress in plants using hyperspectral imaging.

Elements of this review have been published in a Plant Methods journal review paper [13].

2.1 Leaf detection and Segmentation in colour and HS images

Plant images can vary in complexity from a simple leaf with a plain background to multiple overlapping leaves with plants in the background, such as field images. The complexity of data for this research falls in the middle, with overlapping leaves and a plain background, taken in a lab setting.

This project uses RGB data but it also uses hyperspectral data which increases the complexity. Even though hyperspectral images are complex standard computer vision techniques i.e. segmentation approaches can still be used. Therefore the next Section looks at several potential approaches.

2.1.1 Segmentation techniques

The aim of segmentation is to partition the image into sections by grouping similar pixels or superpixels (collection of pixels treated as one that share some characteristics) into sections by labelling the pixels.

There are low level segmentation methods, such as thresholding, which take a value and eliminate everything below or above that value. The thresholding could be based on the intensity or colour. There is also edge detection which looks for a difference between neighbouring pixels to find edges [15] [16]. While there are many situations where these methods would be useful in this case they are not because the data is complex and the thresholding or edge detectors are not good enough to segment leaves. More information is needed in order to find leaves that overlap, and have similar image properties.

The following describe some of the most popular classes of segmentation methods.

2.1.1.1 Clustering

Clustering is a region-based segmentation method and uses features (colour, texture, etc) to group similar pixels (or superpixels) together. K−means and Mean Shift are well known clustering methods widely used in segmentation.

K−means works by calculating distance from points (usually pixels) to the nearest mean and then update the mean to include the new point [17]. If the means have not changed more than a tolerance term ϵ then the algorithm stops. K−means requires an initial K number of points that will become the centroids.

Mean shift was first introduced by Fukunaga [18] and adapted by Comaniciu and Meer [19] for segmentation where the Mean shift uses features in the image such as colour, gradient or texture.

The general algorithm for mean shift is: For each point x_i

1. Find the mean shift vector $m(x_i^t)$
2. Translate density estimation window by $m(x_i^t)$
3. Iterate (1) and (2) until convergence.

Mean shift has been used for object extraction and classification on hyperspectral data [20]. First, matrix factorisation is used to reduce the dimensionality. However it is noted that the kernel density estimation does not work well with

high dimensions and the bandwidth selection is dependent on the data, which means the method needs to be adapted per dataset. The results of Mean shift are very good, and had a higher accuracy than: matrix factorisation, Derivative Morphological Profile (DMP) and Fractal Net Evolution Approach (FNEA) a region merging approach.

2.1.1.2 Graphs

Graph based methods have previously been used to segment plant images (and hyperspectral images). This is because this approach splits sections of the image using statistics and this generally works well. Graphs can partition objects in the image by finding the difference in intensities of the pixels and partitioning the image.

A Graph is represented as G = (V,E) where V is the vertices and E is the edges. Usually the vertices are pixels within the image and the edge is the connection between two vertices. Then the edge is weighted depending on the particular method in use.

The pixels connected to each other are neighbours and there are four or eight neighbours selected depending on the method. The edges are weighted depending on the method, where the weighting could be the similarity of the pixel intensities. The aim is to cut the graph to have the best fit (segmentation) of desired object, where a cut is a separation of the graph into two disjoint sections. There are different graph based methods, some of which are described below.

A simple way of finding the graph cut is to mark the object and mark the background then calculate the probability that each pixel belongs to the class. Then the pixels can have a probability distribution of the likelihood that it is background or object. The minimum cut finds the cut in the cluster of vertices with the maximum flow, where the maximum flow is the most direct cut from the object (source) to the background (sink). Greig et al [21] was the first to use the min cut/max flow in computer vision to minimise energy functions.

The problem is the cuts in min cut tend to happen in the wrong place and not the optimal place along the graph as mentioned in [22]. Instead the authors propose normalised cuts in computer vision. This builds on elements of graph theory however instead of finding the minimum they find the normalised cut. This takes the cut as a fraction of all the edge connections to all of the nodes. This means

isolated nodes will not have a small cut weight and will remain with other isolated nodes.

A relevent application of normalised cuts is used by Gorretta [23]. This focuses on hyperspectral images and uses both the spatial and spectral information as opposed to primarily using the spectral data and ignoring the spatial information. The experiment takes two images; one is a remote sensing image of the city of Pavia in the spectral range 495 to 2000nm over 80 bands, and another is a field visible and near-infrared image of a plant with 3.6nm spectral resolution over the range 400-980nm (160 bands).

The images are split using the normalised cut and the aim is to over segment so that the region merging can correctly merge the regions. The plants segmentation works well however the city segmentation is not fully segmented.

2.1.1.3 Active Contours, implicit and explicit

Active contours move a mathematical line to an object boundary by minimising the energy of a system representing the line. There are control points along the line that are controlled using internal and external forces. Parametric Active Contours are an explicit way of finding the contour; whereas level set approaches find the contour implicitly using partial differential equations, and finding the zero level where the grid interacts with the object.

Kass [24] first introduced active contours in computer vision where the active contours, also known as snakes, have internal forces that are attracted to the edge or line of an object using splines to represent the curve (Eq. 2.1).

$$\int_0^1 E_{snake}(v(s))ds = \int_0^1 E_{int}(v(s)) + E_{image}(v(s)) + E_{con}(v(s))ds \tag{2.1}$$

Where v(s) = x(s),y(s)is the location of the snake, E_{int} controls the shape of the spline/snake, E_{image} are the external image forces (image features such as edges and sharp pixel intensity changes), where the image forces are weighted to attract the snake to the contour. Also E_{con} are extra external constraints on the snake. Snakes can be attracted to an edge of an object from a distance and avoid local minima by using a smoothness constraint to find the object over a noisy image. The weights need to be adjusted depending on the images, also the snakes need to be initialised very close to the desired object otherwise it could locate a different

object or get stuck in local maxima (inner edges, eg. in the case of leaves the veins).

Snakes work on the image plane and directly uses the image features to find the countours. There is a second approach where a function interacts with the image plane in order to find the contours, this is called level sets. Osher [25] first proposed the level set approach, which considers the evolution of a curve over time and implicitly finds the contour of the desired object.

The level set representation is:

$$\phi(p,t) = \pm d \tag{2.2}$$

Where d is the distance from point p to the curve when time $t = 0$. The positive distance represents the outside of the curve, and the negative distance is within the curve. Level sets are partial differential equations, represented as:

$$\phi_t = F(K)|\nabla\phi| \tag{2.3}$$

This follows the Hamilton Jacobi problem. Here K is the curve and $F(K) = F_0 + F_1(K)$. Malladi et al [26] describe the limitations with the original level set and propose narrow bands and reinitialisation. The narrow band level set method focuses on a thin band at the front of the function and therefore instead of computing all the level sets. Reinitialisation is good because when the function is evolving it can become distorted and the function will deviate. Therefore reinitialisation maintains a stable function. The following equations are referred to as the Eularian Hamilton Jacobi formulation.

The function $\phi(x,t)$ evolves over a mesh with uniformly spaced nodes i,j where $\phi = \phi_{ij}$ in the following equations:

$$\frac{\phi^{n+1} - \phi^n}{\triangle t} + (F)(\nabla\phi^n) = 0 \tag{2.4}$$

Where $\triangle t$ is the time step and $\nabla\phi^n$ is an approximate finite difference operator for a spatial derivative.

By rearranging Equation 2.4 the ϕ^{n+1} can be found for the numerical method.

$$\phi^{n+1} = \phi^n - (F\nabla\phi^n))(\triangle t) \tag{2.5}$$

An advantage for this approach is that it can be extended to three dimensions with no significant difference.

Chan and Vese [27] introduce active contours without edges called the Chan-Vese model. This model can find objects or features with contours without gradient, this is not possible with the original snakes or level sets. This approach can also work with noisy images where the object boundaries are less clear. This is helpful for overlapping leaves where the boundary can be visually hard to see.

$$F(c_1, c_2, \phi) = length(\phi) + area(\phi) + \int |f(x) - c_1|^2 + \int |f(x) - c_2|^2 \tag{2.6}$$

Where the piece wise constant intensities of two regions (inside and outside the object), the integral containing c_1 refers to inside the object and c_2 refers to outside the object.

$$length(\phi = 0) = \mu \int \delta(\phi(x))|\nabla\phi(x)|dx \tag{2.7}$$

$$area(\phi \geq 0) = \upsilon \int H(\phi(x))dx \tag{2.8}$$

Extending 2.6 and including 2.7 and 2.8 the energy can be written as:

$$F(c_1, c_2, \phi) = \mu \int \delta(\phi(x))|\nabla\phi(x)|dx + \upsilon \int H(\phi(x))dx$$
$$+\lambda_1 \int |f(x) - c_1|^2 H(\phi(x))dx + \lambda_2 \int |f(x) - c_2|^2 (1 - H(\phi(x)))dx \tag{2.9}$$

Where H, the Heaviside function, is a one-dimensional function:

$$H(\phi) = \begin{cases} 1, & \text{if } \phi > 0 \\ 0, & \text{if } \phi \leq 0 \end{cases} \tag{2.10}$$

2.1.1.4 Spectral unmixing

While spectral unmixing is usually used for distant remote sensing data when the pixels need to be separated, and so it is not so relevant here, however it is briefly discussed below.

Spectral unmixing is the decomposition of a mixed pixel into distinct spectra (end-

members) and the proportion of each endmember within the pixel (abundance). A mixed pixel is when multiple materials are in one pixel. An endmember is referring to raw material such as plant, soil or background. If a pixel is mentioned as pure it means only one endmember is contained in the pixel. [28] list the three steps for unmixing the hyperspectral data:

1. Dimension reduction (optional)
2. Endmember selection
3. Inversion (fractional abundance)

Dimension reduction takes the data and reduces the size by keeping the most important information. Hyperspectral data is usually large, so this process would reduce the size of the data and increase the computation speed. For example each plant that has been imaged for the drought experiment in this thesis has the dimensions (400x1004x812) which corresponds to (x,y,z) axes of the data respectively. There are over 300 million pixels that need to be processed and the amount of pixels would be halved if there were 406 bands instead of 812 bands.

The endmember selection step finds all of the endmembers within the input data matrix. Then the inversion finds the fractional abundance matrix, this is N number of matrices (N = number of endmembers) and each matrix represents the probability of the spectral signatures belonging to that endmember (see Chapter 3.2.2). [29] lists the endmember and inversion techniques available.

Spectral unmixing will either use the linear mixture model or non-linear mixture model depending on if the data is linear or non-linear. [28] explains that the data is linear if there is a linear combination of the endmembers weighted by the fractional area of each endmember in the pixel, i.e. the light interacts with only one component for each distinct endmember. The data is non-linear if the endmembers are mixed at spatially smaller scales than the path length of photons in the mixture. The light interacts with different components as it is scattered.

If the wrong mixture model is selected the proportions of the endmembers within the pixel will be significantly wrong. [28] demonstrates the results of using a linear mixture model with non-linear data.

When the light reflects once off the surface of the endmember in the pixel then the endmember will have a clear signature, which means the fractional abundance per endmember within the pixel will have a linear relationship and the endmembers

abundance will sum to 1. A pixel in the observed matrix can be expressed as:

$$Y_{ij} = \sum_{a=1}^{r} W_{ia} H_{aj} + N \tag{2.11}$$

Where $N \in R^{mxn}$ accounts for additive noise. This can be written as:

$$Y = WH + N \tag{2.12}$$

Where $Y \in R^{mxn}$ is the hyperspectral input data matrix, m is the number of spectral bands and n is the pixels, and $W \in R^{mxr}$ is the endmembers signatures where r is the number of endmembers and $H \in R^{rxn}$ is the abundance matrix containing a fraction for each endmember per pixel. Also each column of H sums to 1 [30].

The linear mixture model has two constraints. The abundance needs to be nonnegative and the abundances need to sum to one. Independent Component Analysis (ICA) has an assumption that the sources are mutually independent. This assumption violates the conditions of the Linear Mixture Model. There are a few constraints applied to ICA to overcome this problem. [31] uses dependency in dependent component analysis by using Dirichlet densities that will satisfy the LMM.

The input matrix is decomposed into two lower rank non-negative matrices containing the endmember signatures and the abundance fractions of each endmember per pixel. The input data matrix is a two dimensional matrix where the pixels spatial location is dismissed and the pixels are converted into a matrix to vector alignment.

$$Y = WH \tag{2.13}$$

Equation 2.13 uses Equation 2.12 and follows the linear mixture model. Matrices W and H are either initialised using random non−negative values or initialised using Vertex component analysis (VCA) for more robust results [32]. Also r (number of endmembers) is either already known, found using Virtual Dimension (VD) [33] or found by trial and error. NMF uses a multiplicative update rule to keep the matrix values above zero and to minimise the matrices. There are several update rules that can be used. Two examples are Lee-Seung (LS) and Alternating Least Squares (ALS). Two minor problems with the update rules are that LS slowly

converge and ALS can reach local minima and diverge ([34])

$$H = H((W^T Y)./(W^T W H)) \tag{2.14}$$

$$W = W((Y H^T)./(W H H^T)) \tag{2.15}$$

$$\sum_{r=1}^{r} H_{rn} = 1 \tag{2.16}$$

$$H_{n+1} = (W_n^T W_n)^{-1}./(W^T Y) \tag{2.17}$$

$$W_{n+1} = Y H_{n+1}^T (H_{n+1}^T H_{n+1})^{-1} \tag{2.18}$$

Equations (2.14), (2.15) and (2.16) are the 'LS' update rule used in [35] and [30]. After the first iteration the estimated matrices are tested for convergence using Equation 2.15. This update continues until the matrices converge or the input iteration number is exceeded.

Equations (2.17) and (2.18) are the Alternating Least Squares update rule, where the matrices may not be non−negative because of the inverse. After Equation 2.17 and 2.18 if W or H values are less than zero then that pixel becomes zero otherwise it keeps the current value [34].

After each iteration of the LS update rule (or another selected rule) the input data matrix (Y) and the estimated data matrices (WH) are compared to check the difference per pixel and if the convergence is below a certain value (i.e. <0.0001) it will stop. The Frobenius norm, equation 2.19, is a popular minimisation equation however other equations can also be used.

$$\frac{1}{n}||Y - WH||_F^2 \tag{2.19}$$

However there may not be a unique solution since WH = (WD)(D-1H) for any nonnegative invertible matrix D which is explained in [36]. This General NMF can be improved using sparseness and smoothness constraints.

2.1.2 Techniques applied to leaf segmentation

Separating the whole plant from the background can be relatively easy if the background is distinct from the plant. It becomes more challenging when the individual leaves need to be segmented. This problem was turned into a competition; since 2014 there has been a leaf segmentation challenge, this was introduced

for the ECCV Computer Vision Problems in Plant Phenotyping workshop where a dataset was created that contained three sets of *Arabidopsis thaliana* and tobacco plants along with the corresponding ground truth segmentation. The three sets of images had different challenging elements; Set 1 had more of a complex background, Set 2 had different varieties of genotypes of *Arabidopsis thaliana* and Set 3 contained higher resolution images.

The first leaf segmentation challenge resulted in several approaches used to count and segment leaves. Three used distance maps and one used template matching, summarised in a collation study [37]. The datasets created for the challenge have been used by researchers since the challenge. Vukadinovic and Polder [38] used the Euclidean distance to create a distance map of the image, then they applied the watershed approach on the map to segment the leaves. Simek and Barnard [39] used a Gaussian process shape model along with a likelihood function to segment the datasets. The shape prior is a combination of the petiole and the blade, the likelihood parameters were trained from the ground truth leaf masks using maximum likelihood. Quan [40] used normalised cuts for part of partitioning the leaves. First colour thresholding is applied, which reduces the amount of nodes needed to be connected. Then the further subdivisions are found using normalised cuts. The focus is on plant modelling using 2D and 3D approaches to find the plant and individual leaves. Also the algorithm relies on a user interface to modify or adjust the segmentation if it is needed, and it does not perform well with occlusions.

Wang et al [41] have adapted a level set approach to overlapping leaves using a two step method. The midrib (main central vein) is first selected (manually placing a line over the vein), then the narrow band evolution region is applied followed by the Local Chan-Vese (LCV) model. Images included two overlapping leaves where the top leaf is segmented.

Machine learning

Machine learning approaches train algorithms using a training dataset, with the aim of analysing and predicting results from new, unseen data. Multilayer perceptron's (MLP) are simple networks (called Artificial Neural Networks) that map input data to an output. An input node connects to the output, and is updated using an activation function and weights that can be optimised to produce the desired output (using training data).

The MLP approach uses a simple architecture consisting of an input, hidden layer(s) and the output. In machine learning a new, more sophisticated approach called deep learning is becoming popular. Deep learning refers to artificial neural networks with a structure that contains a lot of layers. In each layer, neurons are able to implicitly represent features from the data and by doing this, more complex information can be obtained in later layers; image features are automatically determined by the network. One specific example of a deep learning approach is convolutional neural networks (CNN). Whilst ANNs use neuron activation networks as their analogous model, CNNs are based on retinal fields in the vision system. Whatever the approach, deep learning takes longer to train and the architecture can be more complex than traditional neural networks. However, with the added complexity, very impressive classification and recognition rates are achievable. There have been approaches in deep learning over recent years that used plant datasets for leaf counting and segmentation, which have performed well. [42] [43] [44] [45]

2.2 Techniques for early detection of biotic and abiotic stress in plants

The reliable detection and identification of plant disease and plant stress are a current challenge in agriculture [46] [47] . Standard existing methods of detection rely on crop agronomists manually checking the crop for indicator signs that are already visible. Depending on the type of crop and the size of the crop area, which for many commercial crops is often very large, this method of monitoring plant health is both time consuming and demanding. Manual detection also relies on the disease or stress exhibiting clearly visible symptoms, which often manifest at middle to late stages of infection. Identification of the causal agent is through either manual detection or diagnostic tests [48]. Diseases usually start in a small region on the foliage (e.g. *Septoria tritici* blotch (STB) of wheat caused by the fungal pathogen, *Mycosphaerella graminicola*; Apple scab caused by *Venturia inaequalis*), which can be difficult to detect by visual inspection if the crop is large; however, the ability to identify the disease at this early stage would enable an opportunity for early intervention to control, prevent spread of infection, or change crop management practices before the whole crop is infected or damaged. Identifying areas of disease could also lead to targeted application of chemicals. Such precision approaches would result in the reduction of pesticide and herbicide us-

age, with subsequent beneficial impact for the environment, ecosystem services, grower finances and the end consumer. Hence, there is a keen interest in the agricultural and horticultural sector to replace this largely manual process with more automated, objective, and sensitive approaches. Mahlein [49] has discussed the literature on plant disease detection by imaging sensors. This includes RGB, Multi spectral, Hyperspectral, thermal, Chlorophyll Fluorescence and 3D sensors. One conclusion is that RGB and hyperspectral imaging are preferable for identifying specific diseases.

To improve crop management and plant health, several avenues of research are focussing on the identification of the onset of adverse stresses, ideally before ordinarily visible signs are present. Image analysis techniques show much potential here as they represent non-invasive and potentially autonomous approaches to detect biotic and abiotic stress in plants. Image analysis as a research field represents a host of computational techniques which are able to extract information from digital images. From a practical point of view, this means automatic processing of carefully captured images to produce a dataset of desired measurements from the images. The images themselves can come from a variety of sources, from colour digital cameras or smartphones, to more specialist cameras designed to capture a variety of different information in the images. One such technological advance here is hyperspectral imaging, where cameras capture more than the usual three bands of coloured light found in traditional digital imaging. This Section will specifically focus on the subsequent analysis approach known as hyperspectral image analysis. This has recently become financially accessible to a wide variety of users, due to falling technology costs. Analysis approaches are being developed which are enabling the Hyperspectral imaging technologies to be utilised for wider ranging applications. Hyperspectral imaging uses high-fidelity colour reflectance information over a large range of the light spectrum (beyond that of human vision), and thus has potential for identifying subtle changes in plant growth and development.

For plants and vegetation the most useful wavelength ranges to analyse are the visible range combined with near infrared range (Figure 2.1). This wavelength range can capture changes in the leaf pigmentation (400-700nm) and mesophyll cell structure (700-1300nm); however to see water content of the plant changes, extended ranges are needed (1300-2500nm) [50]. Severe dehydration, for example, can affect the leaf mesophyll structure which relates to changes in the near infrared reflectance; however, minor drought stress does not usually have enough of an

effect to be detected [51].

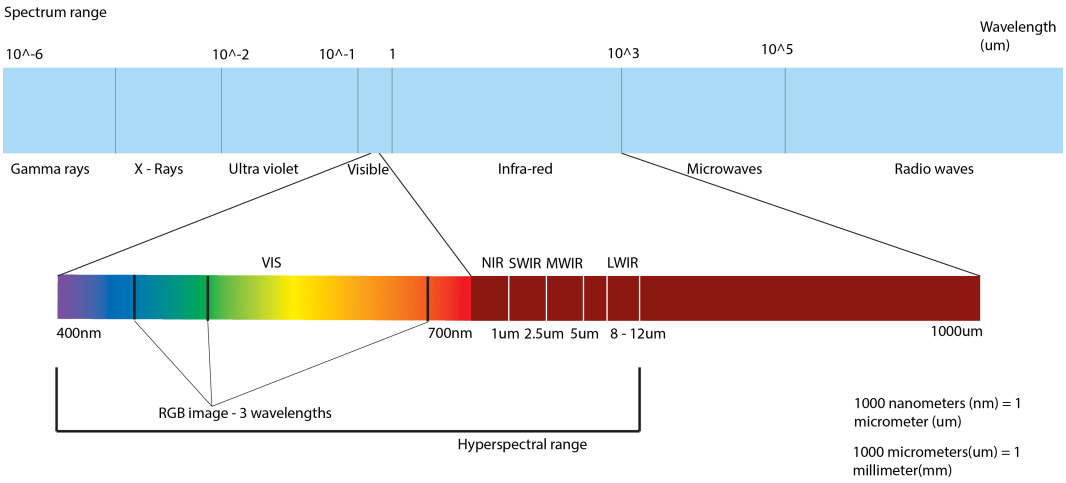

Figure 2.1: Electromagnetic spectrum with the lower bar displaying visible and infra-red light.

2.2.1 Existing Vegetation and Disease Indices

Before hyperspectral imaging devices were readily available, researchers wishing to quantify effects based on colour information have used multispectral imaging, or hyperspectral, point-source devices (such as spectroradiometers which do not produce a spatial image) to acquire colour data. Hyperspectral devices do not in general provide a point-and-click measurement. Instead, much onus is on the user to develop the capture and analysis process. Once acquired, the resulting large numerical datasets must be analysed in order to provide useful information. A sensible and simple way into such large datasets is to consider only a small number of positions in the wavelength range, looking at changes across conditions at pre-determined key points in the spectrum. Using this approach, we can also counter effects of relative light changes by considering ratios of data values. This involves the combination of two or more wavelengths, commonly known as 'indices'.

To interpret the data, a number of such indices have been developed, through either pre-considered biological reasoning (e.g. knowing that a particular wavelength relates to properties in a particular cell structure) or due to limitations in the particular wavelengths available from the capture equipment (e.g. indices which are derived from satellite multispectral remote sensing data may only have

had a limited number of wavelengths available). When applied to plant material, these are known as 'vegetation indices'. Many different vegetation indices exist, and each uses a different set of wavelength measurements for describing physiological attributes of vegetation, looking at either general properties of the plant, or at specific parameters of its growth One of the most popular and widespread metrics is the Normalised Difference Vegetation Index (NDVI) , which is used for measuring the general health status of crops [52] [53]. It is calculated via a simple ratio of near-IR and visible light (see Table 2.1). NDVI has been used for many different purposes, for example, to detect stress caused by the Sunn pest/cereal pest, *Eurygaster integriceps Put. (Hemiptera: Scutelleridae)*, in wheat [8]. Most of the indices are very specific and only work well with the datasets that they were designed for [9]. There are disease-centric studies focused on creating disease indices for detecting and quantifying specific diseases [54], for example, one study used leaf rust disease severity index (LRDSI) with a 87-91% accuracy in detecting the leaf rust (*Puccinia triticina*) in Wheat [55].

Table 2.1: A selection of Vegetation indices

VI	Formula	Reference	Information
Normalised difference vegetation index (NDVI)	(RNIR - RRED)/ (RNIR+RRED) RRED 680, RNIR 800	[56]	Range: -1 to 1 GH: 0.2-0.8 Broadband
Red Edge NDVI	(R750-R705)/ (R750+R705)	[56]	Range: -1 to 1 GH: 0.2 to 0.9 Narrowband (hyperspectral data)
Simple ratio index(SRI)	RNIR/RRED RRED 680 RNIR 800	[55]	Range: 0 to 30 GH: 2-8 Broadband
Photochemical Reflectance Index (PRI)	(R531-R570)/ (R731+R750)	[56] [57]	Range: -1 to 1 GH: -0.2 to 0.2 Vegetation health prior to senescence
Plant Senescence Reflectance Index (PSRI)	(R680-R500)/ R750	[56]	Range: -1 to 1 GH: -0.1 to 0.2 PSRI used for canopy stress, onset of senescence, fruit ripening
Normalised Phaeophytinization Index (NPQI)	(R415- R435)/(R415+R435)	[58]	Chlorophyll degradation 0.56 to 1.41* *unacidified and acidified solutions
Structure Independent Pigment Index (SIPI)	(R800-R445)/ (R800+R680)	[56] [57] [59]	Range: 0 to 2 GH: 0.8 to 1.8 Good with canopy variety
Leaf Rust disease Severity Index (LRDSI)	6.9* (R605/R455)- 1.2	[55]	Accuracy of 89% in study may vary with other data

Another commonly-used approach is to detect changes in the sudden increase in reflectance at the red/near-infrared border. This 'red edge' position is a narrow section in the electromagnetic spectrum (690-740nm) where the visible spectrum ends and the near infrared starts (Figure 2.2). This section has a large derivative within leaves due to the amount of reflectance in the near infrared that the plants emit. Cho[60] describes a number of different algorithms that extract or detect the red edge. A disease index based on the red edge position has been used to detect powdery mildew in wheat (*Blumeria graminis f. sp. Tritici*), however it

was not found to be as accurate as Partial Least Squares Regression (PLSR), a technique that uses a statistical approach [61].

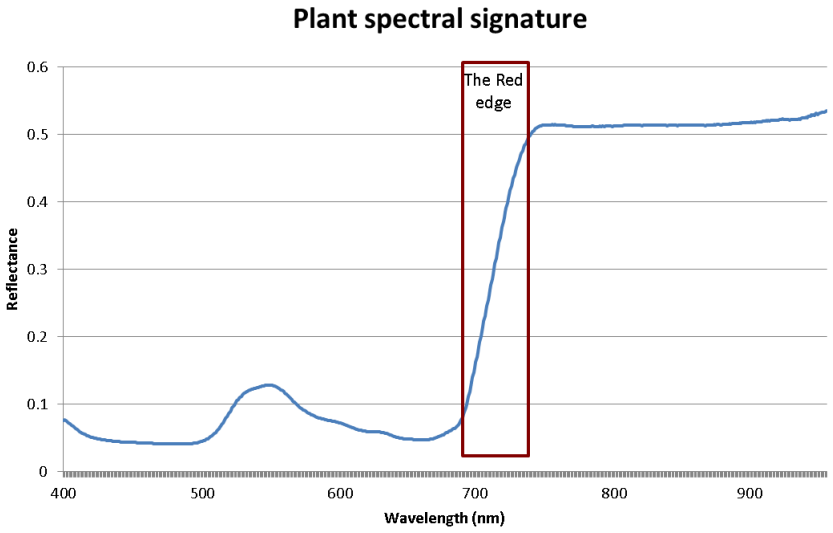

Figure 2.2: A typical healthy vegetation spectra (400-1000nm) with the red edge section highlighted in red (690-740nm).

2.2.2 Applications for the detection and classification of healthy and diseased stressed plants

A variety of techniques are used specifically for the detection of biotic stress in plants. Classification techniques, that is, those that separate the data into healthy and diseased categories for example, can be divided into two types: those that use the entire spectrum response, and those that focus on a number of key wavelengths in the spectrum. Furthermore, disease classification is discussed with regards to the identification of multiple diseases and detection of a specific disease.

2.2.2.1 Classification using a subset of selected wavelengths

In this Section we consider classification approaches that truly rely on multispectral data, that is, sub sampling at particular wavelengths in the full spectrum. With true multi-spectral data specific wavelengths can be manually or automatically chosen from anywhere in the captured range, as the full range is available in a hyperspectral dataset.

Analysis from Section 2.2.1 typically used indices to calculate representative values using discrete wavelengths at various positions in the spectrum. One such

study involving a wheat field experiment used Normalised Difference Vegetation Index (NDVI) response to eliminate everything except the leaves from the dataset, followed by a statistical approach called an ANCOVA (which measures statistical covariance) to identify selected wavelength bands, and then Quadratic Discriminant Analysis (QDA) to classify the spectra between healthy and diseased leaves (yellow rust)[62]. This is representative of a typical workflow in hyperspectral analysis: isolate (or segment) the parts of the image of interest, then use a mathematical technique to identify regions of the spectra likely to have predictive power, and finally use those spatial and spectral regions to learn a classification approach. Using QDA, the overall accuracy reached 92% with 4 wavebands [62].

An example of a Multilayer perceptron(MLP) is described by Moshou [63], who aimed to detect yellow rust in field-grown wheat using a spectrograph with the range 460-900 nm and a 20 nm spectral resolution. The spectrograph captured the images in the field using a handheld system. Then four significant wavelengths were selected. The first two wavelengths were selected using 'variable selection' which involved comparing the wavelengths using stepwise discriminant analysis and using the F-test. The second pair of wavelengths uses the NDVI wavelengths. The neural network used by Moshou is a simple architecture with four inputs, one hidden layer consisting of ten neurons and two outputs (healthy and diseased). The architecture is determined by the number of inputs, a selected amount of hidden neurons and the amount of outputs required. Trial and error can be used to determine a suitable architecture. Moshou tried different quantities of neurons and selected the most efficient. The classification accuracy reached using this approach was 98.9% for the healthy plants and 99.4% for the diseased plants.

Deep learning has been applied to the problem of plant disease detection. Mohanty [64] used CNNs to detect 26 diseases over 14 crop species. A dataset consisting of 54,306 colour images were used, 80% for training and 20% for testing on AlexNet and GoogLeNet (two popular versions of pretrained CNN's). The accuracy was 97.82% for AlexNet and 98.36% for GoogLeNet using colour images. They selected individual leaves with a homogenous background. If the network is tested on images under different conditions from the trained images the accuracy is 31.4% [64]. Sladojevic also used CNN's to detect 13 diseases across various crop plants, including Apple (powdery mildew, rust), pear (leaf spot), grapevine (wilt, mites, powdery mildew, downey mildew) using 30000 images with an accuracy of 96.3% using CaffeNet [65].

There are currently very few complete studies applying deep learning to hyper-

spectral data, though this is an active research area. There are several challenges that need to be addressed in order to use hyperspectral data for deep learning. The size of the hyperspectral data including the amount of wavelengths would require a lot of processing time and power it would ideally require a graphics processing unit (GPU). The amount of hyperspectral wavelengths would most likely include noise from specific wavelengths. Also there needs to be a sufficient amount of data for the training/testing process along with labelled data. There is also the possibility that the error will be higher than alternative approaches.

Other non-deep learning approaches include using Fishers Linear Discriminant Analysis with remote sensing data to detect yellow rust and powdery mildew for a wheat crop with an overall accuracy of 93% with selected wavelength ranges (531nm, 570-654nm, 685-717nm) that are significant for detecting differences between powdery mildew and yellow rust diseases in these spectral reflectance ranges, resulting from an independent t-test [66].

Sometimes data analysis approaches are combined with simple image processing steps in order to add feature discrimination. A family of image processing techniques called morphological operators can be used to clean up binary (black and white) images. One such technique is called erosion, whereby the foreground of an object is shrunk by turning boundary pixels into background pixels. The opposite technique is called 'dilation' and has the effect of enlarging the foreground objects boundary. They can be used together to fill in holes, or remove speckle noise (depending on the order used) in binary labelled data. One approach using this method is a study on cucumber leaf data has been used to analyse a different type of mildew; downy mildew (*Pseudoperonospora cubensis*). First Principle Component analysis (PCA) is applied to reduce the size of the data and a binary image is produced, and then Erosion and Dilation are used in a second step to enhance the disease features. The accuracy is 90% however only 20 samples were used (10 healthy and 10 infected) [67]. This method is unlikely to work as well on other hyperspectral images to detect diseases unless the leaf data is similar and even then the results are uncertain.

Hyperspectral imaging can also be combined with microscopy to capture images at a higher resolution. Barley with different genotypes has been studied at the microscopic level to see if spectral differences could be identified between the genotypes. Barley leaves were also analysed from both healthy and diseased plants, which had been inoculated with Powdery Mildew (*Blumeria graminis*). Results showed there was a difference over time between the healthy and inoculated leaves, except for

those varieties containing the mildew locus o (mlo) gene, which provides plant resistance to Blumeria graminis. In this study, the spectral range was reduced to 420-830 nm due to the noise, then normalised and smoothed with Savitzky-Golay filter, and then Simplex Volume Maximisation (SiVM) is used to find the extreme spectra followed by Dirichlet aggregation regression for the leaf trace [68].

2.2.2.2 Classification using full spectrum data

Classification approaches aim to divide the data into a number of distinct classes. One such approach is Quadratic Discriminant Analysis (QDA), which classifies by using a covariance matrix, which compares classes. The QDA method was used in a study with Avocado plants, to examine the fungal disease Laurel wilt (*Raffaelea lauricola*), using plants located both in the field and glasshouse. The QDA classification accuracy was 94% [69]. It is possible of course to use alternative methods at each stage of the analysis pipeline. For example, rather than use QDA, a Decision Tree approach (a machine learning technique) has been used and reached 95% accuracy [69]. Choosing the correct approach for the data, as well as ensuring sufficient dataset size and quality, is key.

MLP's requires prior knowledge (training data) therefore if the 'disease spectra' is unknown then this technique will be unsuitable.

A third classification approach is to look at the spectral signatures by using derivatives; this is when the underlying pattern or change in data is analysed. Second order (and above) derivatives are usually insensitive to changes in the illumination [70]; however they are sensitive to noise which hyperspectral data typically suffers from, therefore 'smoothing' needs to be applied before using derivatives. Smoothing is a process that reduces the difference between individual pixel intensities and neighbouring pixels using forms of averaging to create a smoother signal. Two smoothing examples are Savitsky-Golay and Mean filters. Savitszky-Golay proposed a method for smoothing noisy data by fitting local polynomials to a sub set of the input data then evaluating the polynomial at a single point to smooth the signal [71]. The Mean filter replaces pixel values with the mean of the neighbouring pixels, this helps to reduce noisy pixels.

Huang [72] tried to detect Sclerotinia rot disease in Celery crops by using Partial Least Squares Regression (PLSR) with derivatives of first and second order. Partial least squares regression reduces the predictors to a small set of components. This technique is useful when the predictors are collinear/highly correlated and you wish to reduce the set to a smaller set of key values. The classification

accuracy for PLSR with the raw spectra is 88.92%, PLSR with Savitzky-Golay first derivative is 88.18% and PLSR with second order derivative is 86.38%. The accuracies are similar, with the second order derivative performing slightly worse. Yuan used PLSR on Fishers Linear Discriminant Analysis (FLDA) to detect pest and disease in wheat. It produced a 60% accuracy for aphid damage and a 92% accuracy for Yellow rust disease [73]. In another study, Zhang used FLDA to detect powdery mildew in wheat (using heavily damaged leaves) with over 90% accuracy [74].

2.2.2.3 Disease identification

As well as detecting the presence of disease, another avenue of research is to distinguish between different diseases to identify specific pathogens. One such approach is spectral information divergence classification. This method compares the divergence between the observed spectra and a reference spectra (a library of spectra, or average spectra of interest from the data), where the smaller the divergence value then the more similar the spectra are, and if they are larger than a set threshold then they are not classified as the reference spectra [75]. Spectral information divergence was used to detect canker legions on citrus fruit (grapefruits) where the spectral range of the data was 450-930 nm with 92 bands and 5.2 nm spectral resolution. Before analysing the data, a pre-processing step is applied by combining neighbouring pixels to reduce the size by half. Cankerous grapefruits were compared with normal grapefruit and also with grapefruit showing other disease or damage symptoms including: greasy spot, insect damage, melanose, scab and wind scar; this method resulted in 95.2% classification accuracy [76].

2.2.3 Quantifying severity of disease

Along with detecting and classifying disease, we may wish to record the effective amount of disease, or its severity. This approach does run into some particular challenges. The amount of leaf damage and coverage from the disease can affect the accuracy of the leaves being classified as healthy or diseased. Extreme disease damage can affect the appearance of leaves so detrimentally that they may not be counted as plant material at all. Still, there are a number of methods for estimating severity, and we present a selection of approaches below.

Spectral Angle Mapper (SAM) approaches match the pixel spectra to reference spectra to classify the pixels by calculating the angle between the spectra which are treated as n-dimensional vectors in space [77]. This technique has been widely

used with moderate success to classify hyperspectral data, including plant diseases. Yuhas studied the severity of Fusarium head blight disease for wheat before harvesting. The hyperspectral data was in the range 400-1000 nm with a spectral resolution of 2.5 nm. SAM was used to detect the amount of disease with a classification accuracy of 87%. Two experiments with wheat plants were carried out, one in a glasshouse and one in field. The plants were imaged over their developmental stages from inoculated to established infection. Yuhas determined that just after infection, the healthy and infected plants were not distinguishable because the infection had not yet established. However, when the hyperspectral data were examined during the ripening stage, the wheat pigment composition changes, and the healthy plants then appear as diseased plants [78].

Mahlein used the same technique to analyse sugar beet diseases specifically Cerospora leaf spot, powdery mildew and leaf rust. The range is 400 -1000 nm with 2.8 nm spectral resolution and 0.19 mm spatial resolution. The plants were analysed over a time period (>20 days) to monitor the different stages of each disease, and the leaves were classified as healthy or diseased. Cerospora leaf spot classification accuracy varied depending on the severity of the disease (89.01% to 98.90%), powdery mildew accuracy varied between 90.18% and 97.23%, and sugar beet rust reached 61.70%, with no classification before day 20 using SAM [79].

Rumpf et al (2010) used the same dataset as Mahlein but with different analysis approaches; Decision Trees (DT), Artificial Neural Networks (ANN) and Support Vector Machine (SVM). All approaches require prior knowledge, however once trained have proven to be efficient. For example, with Cerospora leaf spot the accuracy for SVM is 97% (compared to DT- 95% and ANN- 96%); for Sugar beet rust the accuracy is 93% (DT- 92%, ANN- 95%); and for Powdery mildew the accuracy is 93% (DT− 86%, ANN− 91%). Measuring the severity with leaf area coverage after the disease has covered 1-2% of the leaf the accuracy is 62-68% and for more than 10% leaf coverage the accuracy is almost 100% [80]. This demonstrates that it is possible to use a variety of analysis methods on the same set of hyperspectral data to elucidate different insights and achieve different levels of accuracy - choice of technique is important.

2.2.4 Detection of early stage stress symptoms

The ultimate goal of such detection systems is to identify the disease with a minimum of physical changes to the plant. Identifying diseases or abiotic problems as

early as possible has obvious benefits. By using hyperspectral technology in combination with appropriate analysis methods, we can realistically hope to identify stress symptoms before a human observer.

Drought can be a significant problem for many crops [81], particularly as some plant species or varieties do not visibly indicate this stress for a period of time, and by this time, the potential yield or quality of the crop may have decreased because normal plant developmental processes have been affected through the stress response. The definition of 'drought' can also vary from a little water deprivation to complete deprivation. Studies discussed in this Section have detected the onset of drought before Vegetation Indices' detected the drought and also days before visible signs appeared.

One technique in particular which become popular for early detection of drought stress is SiVM, which is a data clustering technique [82]. This technique selects spectral signatures that are samples of healthy and stressed plants, and then clusters the data using these classes. When the signatures become similar to a prelearned sample signature then it is classified as such.

Romer [83] studied drought stress in a barley experiment contained in a rainout shelter and a corn experiment grown in field. The technique used to detect the stress was again Simplex Volume Maximisation, which is unsupervised. The spectrum range was 400-900nm, with 4nm spectral resolution. During pre-processing some wavelengths are removed due to noise (<470 and >750nm). This is a common occurrence with hyperspectral data due to unstable illumination at the end of the spectrum range, and is especially common with lab-based light sources which may not generate much light in these regions of the spectrum. To reduce the size of the data and to remove the background, a K-means clustering method was used to separate the data into a selected number of groups using mean colour. SiVM is then compared to four well known vegetation indices'- NDVI, Photochemical Reflectance Index (PRI), Red Edge Inflection Point (REIP) and Carotenoid Reflectance Index (CRI). For the Barley data, reduced partial water stress was detected four days earlier with SiVM (day 9) than Vegetation Indices (day 13). For the plants with no water/ complete drought conditions the Vegetation Indices 'detected the stress on day 8, one day faster than SiVM but they failed to detect the stress for days 9 and 10.

Behmann also analysed drought stress in barley using Support Vector Machine (SVM). This algorithm is supervised and requires labelled training data, which in this case is labelled as drought or healthy. The data is pre-processed with K-means

to reduce the size of the dataset before analysis with SVM. Using this approach, Behmann detected drought stress on day 6, with NDVI detecting a difference on day 16 [84].

Drought stress in wheat has been analysed by two combined techniques to try and improve detection rates. Moshou [85] uses least squares support vector machine (LSSVM) to try and detect drought stress. Wheat plants were studied in a glasshouse and both spectral reflectance and fluorescence were analysed. LSSVM needed to be trained, and 846 data samples were used for this training, whilst 302 data samples were used for the testing stage. For some techniques the size of the dataset and/or number of wavelengths will determine the time taken to analyse the data due to computation time. Therefore, Moshou used six wavelengths - 503, 545, 566, 608, 860 and 881nm. The LSSVM attained 76.3% accuracy for stress leaves and 86.6% accuracy for healthy leaves. However, the study stated that by using a fusion LSSVM model combining spectral and florescence features, the overall accuracy was greater than 99%. Fluorescence is the measure of chlorophyll fluorescence in the leaf to determine physiological changes.

According to Kersting [86] many of these techniques are difficult to use for non-machine learning or data mining experts because the hyperspectral data needs pre-processing or adapting (i.e. finding the leaves or using select wavelengths). In addition, the other techniques apart from [83] do not analyse lots of plants over several days. This is an important factor to consider for plant phenotyping when there is a lot of data to analyse. Kersting claims to have the first Artificial Intelligence technique for drought stress prediction using hyperspectral data. A novel approach is developed which includes a predictive technique for drought that does not adapt the data or reduce the size. Kersting demonstrates the approach in a Barley drought experiment with data collected over a five-week period. The technique used is called Dirichlet aggregation regression (DAR) and it is based on matrix factorisation. First Simplex Volume Maximisation is used to find 50 spectral signatures from the data and classify them. Then, latent Dirichlet aggregation values are estimated before using a Gaussian process over the values to find the drought levels per plant and per time point. Finally, the process predicts the drought-affected plants before there are visible signs. Based on a five-week barley experiment, prediction of drought occurred 1.5 weeks before visible signs appeared. A comparison of runtimes between SiVM and DAR was assessed and resulted in a runtime of 30 minutes for parallelized SiVM, versus only several minutes using the DAR model. This demonstrates that developing custom analysis techniques

can outperform (either in computation time, required assumptions, ease of use, or final accuracy) the direct application of existing approaches.

There has been a significant increase in scientific literature in recent years focusing on detecting stress in plants using hyperspectral image analysis. Plant disease detection is a major activity in the management of crop plants in both agriculture and horticulture. Hyperspectral imaging is a non-invasive process where the plants are scanned to collect high-resolution data. The technology is becoming more popular since the falling costs of camera production have enabled researchers and developers greater access to this technology.

2.3 Summary

Techniques to segment objects in images have been described in Section 2.1 specifically related to segmenting individual leaves. Then the CVPPP challenge and work related to the datasets have been discussed, as it involves counting and segmenting leaves.

Chapter 3 will present an initial look at the data. Non negative matrix factorisation will be used to explore the hyperspectral data in the raw format and treating the data as unsorted data.

Distance maps and active contours have been covered in this Chapter, level sets are used in Chapter 4 for the leaf segmentation model. Active contours are a promising approach for segmenting leaves and finding the contours and information within the image, especially the Chan-Vese model that can find edges when there is noise or the edges are hard to distinguish. Shape priors have also been found to be useful when the objects are hard to detect ie overlapping similar shapes. Level sets are more reliable than snakes due to the function interacting with the image plane implicitly rather than explicitly, also snakes work in a local range whereas the level sets can be global.

Section 2.2 includes various techniques available to analyse the data to detect biotic and abiotic stress in plants, examples of which have been discussed, with a focus on the classification of healthy and diseased plants, the severity of disease and early detection of stress symptoms. This is explored in Chapter 7 to analyse the data.

Chapter 3

Data collection and initial manual analysis

This Chapter comprises of two main Sections. The first Section describes collecting and calibrating the data. The second Section looks at an initial inspection of the data by manually selecting leaves and examining a drought dataset over the time series from healthy to visibly stressed. This provides an initial insight into the data analysis challenge, and a baseline for the automated data extraction and analysis covered in later thesis Chapters.

A quick overview of the data collected is in Table 3.1

Table 3.1: Data collected

Stress	Start date	Days	Plants	Instances	Time steps	Location
Spider mite	15.06.15	1	Elsanta	10	2	NIAB EMR
			Fenella	10		
Powdery mildew	22.06.15	16	Elsanta	13	9	NIAB EMR
			Fenella	12		
Drought (E)	19.07.15	5	Fenella	21	5	NIAB EMR
Drought (N)	13.08.15	6	Sonata	6	4	Nottingham
			BSP14	12	4	
			Malling Centenary	8	3	
			Hawaii4	12	3	

Drought(E) is the drought dataset collected at NIAB EMR and Drought(N) is collected at Nottingham.

Problems that occurred with data capture, in summary:

- Spider mite - Biological problem: The over saturation of spider mites during the infection process resulted in visible signs appearing at the next imaging time point.
- Mildew and Drought(E) - Technical problem with the HS camera: The images included a visual shift in the wavelengths. This created a 3D effect in the images due to the misalignment in the camera sensors. The spectral signatures were affected.

Due to these reasons only the Drought (Nottingham) dataset is used later in this thesis, however the process of capturing these datasets are discussed in more detail below.

3.1 Part 1: Data capture

3.1.1 Imaging set up

During this project the data was collected at two locations: NIAB East Malling Research (EMR) in Kent, and The University of Nottingham − Sutton Bonington campus. The hyperspectral imaging system specification varies in each location, see Table 3.2.

Table 3.2: HSI details for the two cameras used to collect data in this Chapter.

Location	UoN Sutton Bonington	NIAB East Malling Research
Make	Headwall Photonics	IMEC Ltd (Gent, Belgium)
Type	Pushbroom	Pushbroom
Range	400-1000nm	600-975nm
Wavelength resolution	0.74nm	4nm
Image size	400x1004x812	2048xlength scan

The system set ups are both based in a laboratory with the camera attached to a metal frame for support. The Nottingham system involves moving the camera, whereas the EMR camera is stationary and the object moves on a platform. Therefore, both are pushbroom systems, but implemented in two diferent ways.

The camera at the University of Nottingham captures visible and near infrared (VNIR) wavelengths, which has the range 400 –1000nm and is comprised of contiguous bands where each band is 0.74nm wide with a total of 812 bands. The Spidermite, Mildew and Drought(E) experiments were collected at NIAB EMR using the camera in Figure 3.1b. Drought(N) dataset was captured in Nottingham using the camera in Figure 3.1a.

Hyperspectral data capture and software

Hyperspectral data is large in size, especially when multiple plants are imaged for several days. A scan of a single plant could easily be around a gigabyte in size. If the whole spectrum range is analysed then the process will take considerably longer than selecting several wavelengths to analyse. However, there is a lot of information contained in the data, which could be valuable. If the camera collects 800 spectral bands, are all 800 needed or would binning into 400 or 200 etc. bands be sufficient?

Storing fewer spectral bands results in smaller file sizes, and reduces the complexity of the data analysis, at the expense of throwing away potentially important colour properties. Polder et al. [87] explores the calibration and characterisation of spectrographs captured using three system set ups. The experiments look at the different types of noise and signal-to-noise ratio. The experiments also determined that to an extent binning can occur without loss of information by calculating the resolution, the spectral range and the amount of pixels.

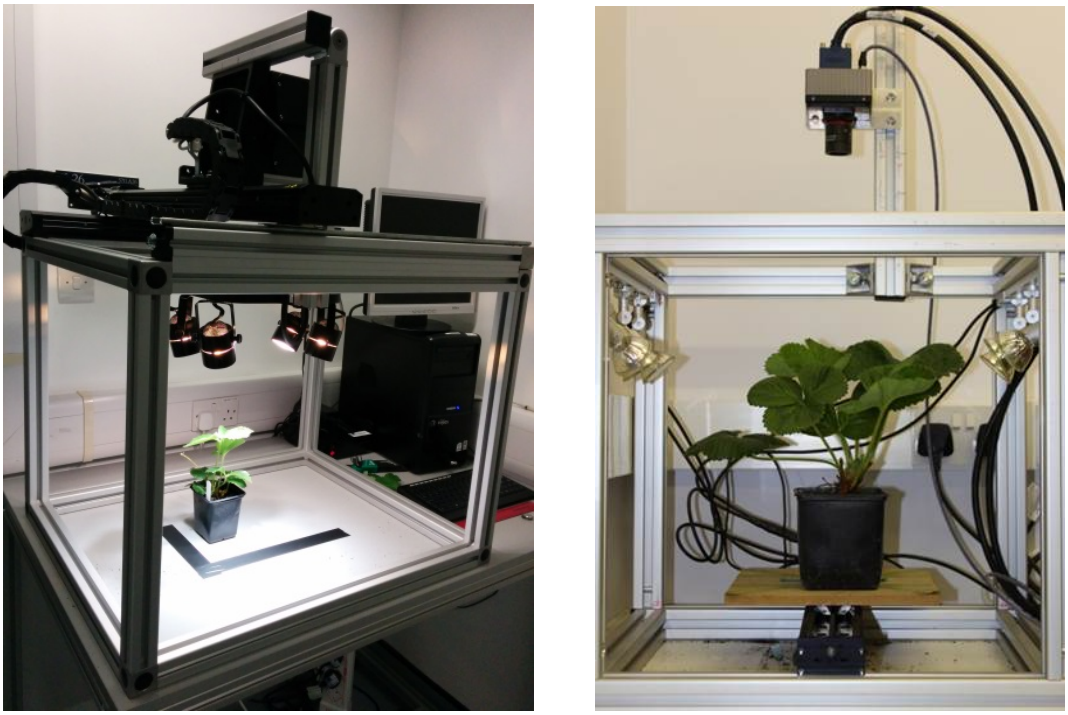

(a) UoN camera set up **(b)** NIAB EMR camera set up

Figure 3.1: (a) the Nottingham camera system and (b) the EMR camera system. The square frame holds the camera and the cables so the camera can move and capture the plant (a) or the plant is moved (b). The Nottingham system has the capability to move the camera in the x-and y- axis, but for imaging the strawberry plants the camera only needs to move along the -y axis.

3.1.2 Selection of plant varieties and type of stress/disease

There are a lot of strawberry varieties that could have been used in this project. The varieties selected have a good variation of traditional and modern types of strawberry plants, also representing a good variety of architecture and stress susceptability. Six varieties were selected: Elsanta, Fenella, Sonatta, BSP 14, Malling Centenary and Hawaii-4. Elsanta and Fenella were selected because they have some resistance to powdery mildew. These were grown from cold store to minimise the chance that the plants were exposed to more than the intended stress or disease. Also this limited the height of the strawberry plants; if they are too large this can cause problems for imaging. Figure 3.1b displays a medium sized Elsanta plant on the imaging board where the leaves are at the same height as the spot lights which means the plant will not be illuminated and the image will be dark. Sonatta, BSP 14, Malling Centenary and Hawaii-4 were existing mature plants with an increased chance that the plants were exposed to other stress and

diseases.

Three stresses were investigated before deciding which stress would be selected for further study in this thesis, three initial experiments were conducted to investigate the development of these stresses.

Two biotic stresses were selected; Powdery Mildew and Spider Mites. Powdery mildew can devastate crops because it is spread by wind dispersal from an infected plant nearby. The spores land on the leaves and the infection takes 14 days (10 if the leaves are directly under the infected plant) to show the first signs, which are curling leaves, small white patches and spots underneath the leaves. Powdery mildew is specific to a plant species which means strawberry powdery mildew will only affect other strawberry plants; however it can affect the whole crop through wind dispersal and by the time the visible signs appear the plants productivity is reduced. Due to the 14 day window there is a lot of time that could be saved by noticing the stress early.

Spider mites are incredibly hard to see by eye. They are tiny (0.4mm), translucent and hide under the leaves where they cause damage. A solution to spider mites is larger pests that only eat the spider mites and not the plants; however there need to be signs of spider mites before application. If there are a few spider mites (<10) then the period from infection to visible signs can be up to several days.

The third abiotic stress selected was drought. This is a broad term for restricting the plants water intake. It can be anything from an increase in water deprivation to complete water withdrawal. For this project the complete withdrawal option was taken.

The experiments lasted until the first visible signs of the relevant stress appeared. There is a protocol for the drought process below and the powdery mildew and Spider mite protocols are in Appendix B which explains each experiment including the process, equipment needed, how the plants were stored and the imaging process. The data collected is listed in Table 3.4.

3.1.3 Protocol for imaging the strawberry plants during stress application

Protocols were developed at NIAB EMR to image strawberry plants with a hyperspectral imaging camera. A protocol ensures appropriate procedures and precautions are taken, including how to store the plants, how to inoculate or infect them

and how to contain the disease/pest while imaging. The drought(E) protocol is included below while the powdery mildew and spider mite protocols are in the Appendix.

Protocol Strawberry Plants - drought experiment

Plant Material Preparation

Dormant strawberry plants were purchased from strawberry propagators and kept in a cold store at 2°C until required. The dormant strawberry plants were taken out of the cold store the night before potting and left at +2°C to ensure plants were not damaged due to a sudden change in temperature. The following day, the strawberry plants were potted into FP9 pots using standard compost and grown for 2-3 weeks until they have 2-3 leaves. When plants have 2-3 leaves, a thorough inspection of the leaves was undertaken and any marks/signs on the leaves of any other diseases/pests were noted.

Equipment: FP9 pots, Standard compost, Cold store plants, Glasshouse for the plants to grow in.

Location for the plants during the experiment: The plants were kept in a heated glasshouse.

Drought preparation and process

Stop watering the plant on day zero.

Equipment: Box with lid and box, 14 Fenella plants (7 drought, 7 control)

Imaging Process

The strawberry plants were imaged on day 0. They were then imaged every day until there were visible signs on all drought induced plants. To transport the plants from the glasshouse for imaging, a box with a lid was used. Also after each imaging session the plants were placed back into random positions.

Imaging method at EMR

1. Clean the area around the camera with 70% ethanol
2. Turn on the three switches at the plugs and turn the PC on.
3. Open the HSI software and select 'start pipeline'
4. Switch the lights on at the voltage box, bottom left and then bottom right.
5. Move the board using the arrow keys on the screen if it is needed.
6. Turn the main room light off.

7. Place the white balance on the board.
8. Select the white balance calibration button
9. Cover lens with hand and select dark (white) calibration button
10. Cover lens with hand and select dark (object) calibration button
11. Place the plant on the board under the lens.
12. Make sure the lens is in focus.
13. Move the board to the desired start position using the arrow keys on the screen

    (Repeat steps 14-20 for each plant)
14. Take a digital image of the plant
15. Place the plant on the board
16. Click the right arrow key on screen to take the board to the start position
17. (select right arrow key to let the plants speed be tested, the image on the screen needs to match on the left and right sides)
18. Change the number of the speed slightly and repeat the last three steps until the image matches on the screen.
19. Name the file a relevant name ie. Day0_Elsanta1_Cal
20. Select frame arrow right key on screen to finally capture the images.

Equipment:

- Camera and Frame set up
- Lighting and set up
- Board
- PC and software - HSI
- Plants and box (to carry the plants in)

Recording information during imaging (note): plant variety, number, date, time and time after treatment, Speed, start and stop position, lighting, exposure rate and frame rate. Any visible marks/signs (including before infection to discount it). General health.

3.1.4 Stress experiments

All three types of stress were imaged to identify the most suitable stress with which to develop the methods for the rest of the thesis. Details and results are presented below.

Powdery Mildew

Mildew can typically take 14 days from infection before visible signs appear. The signs of mildew are curling of the leaves, white spores on the underside of the curled leaves and small dark patches form on the underside. This experiment took place over 16 days in June/July 2015 at EMR.

The plants used were 14 Fenella (7 control and 7 infected) and 11 Elsanta (5 control and 6 infected). The sixth plant in the Elsanta set was used for destructive sampling by removing leaves for imaging.

Due to the length of the experiment, the plan was to image every other day at the same time each day. This started with day 0 (before infection). The plants would be placed into controlled environment cabinets to control as many conditions as possible (Figure 3.2a). After imaging, the plants would be placed back in the Mildew cabinet at random positions to prevent bias in the data. The two Mildew infected host plants were raised above the experimental plants for imaging to enable the mildew spores to naturally drop/float onto them to provide suitable inoculum for good infection (3.2a).

The temperature was set at 19°C for the day and 17°C for the night, and the humidity was set at 70%. The day/night cycle was 15hrs day/9hrs night. The plants were watered every other day.

(a) The Mildew infected plants

(b) The control plants

Figure 3.2: (a) The Mildew infected cabinet set up (b) The control plants cabinet set up.

Spider Mite

The Spider Mite infection lasted for just 24hrs until visible signs appeared. This was due to over saturation of spider mites. The inoculation process involved selecting a small quantity of spider mites on a section of leaf from an infected plant and transferring the leaf to the strawberry plants. It was hard to determine the exact amount of spidermites under the microscope due to the size of the young mites and eggs.

Ten Fenella (five control and five infected) and ten Elsanta (five control and five infected) plants were used. Before they were infected the plants were imaged for a control dataset, and then transported to a glasshouse containing spider mite infected plants. To contain the spider mites on each plant the plants were placed on an overturned plant saucer and placed on a tray filled with water (3.3). Spider mites do not like the water and a moat kept the spider mites contained on the plants.

(a) The Spidermite infected plants

(b) The control plants

Figure 3.3: (a) The Spider mite set up in a controlled glasshouse, (b) The control plants set up.

Spider mite infection involved taking a leaf from a previously infected host plant and placing it at the base of the stem of the plant to be inoculated. Using this method enables the spider mites to crawl onto the leaves naturally rather than being placed on the leaf (Figure 3.4).

(a) A Spidermite crawling up the stem

(b) Leaf placement for infection

Figure 3.4: (a) An example of a Spider mite set in the process of naturally moving to the leaf, (b) The placement of the leaf for infection.

The leaves used for inoculation had many spider mites (50+), therefore by the next day there were already visible signs of infection. The plants were imaged and then the data collection ended because visible signs of infection were already present. The aim of this experiment was to have a slow progression for a time series dataset over several days; however, within 24 hours easily visible signs were already present.

Drought

Two drought datasets were captured. The first drought dataset was collected at EMR NIAB over 10 days in July 2015. 21 Strawberry plants cv. 'Fenella', were placed into a heated glasshouse (Figure 3.5). The experiment contained 7 control plants, 7 drought plants and 7 over-watered plants. The plants were imaged daily until day 4 when the drought plants showed visible drooping/weakness by this point.

Figure 3.5: The set up and location for the Fenella plants during the drought stress and imaging process at EMR

At Nottingham, a second drought experiment took place using the cultivars Hawaii4(12), BSP14(12), Malling Centenary(8) and Sonata(6) during two weeks in August (Figure 3.6). These plants were selected because they supply a variation within the strawberry plant varieties. Sonata and Malling Centenary are commercial plants, BSP14 is visually different and Hawaii4 is a wild strawberry cultivar. The plants were evenly divided into control or drought plants and they were imaged daily for 5-6 days, depending on the plant variety and its sensitivity to drought. The data was collected until visible signs of stress appeared on all of the plants, therefore some plants were imaged for longer. Sonata was imaged for 6 days, there was also a day 9 collection where visible signs were present, and two plants (one control and one drought) were imaged on day 10. BSP14 was imaged for 6 days, Malling Centenary was imaged for 5 days and Hawaii4 was imaged for 6 days. The imaging details are listed in Table 3.4.

Figure 3.6: The set up and location for the plants during the drought stress and imaging process at Nottingham

3.1.5 Data collected: summary

Table 3.3: Table of the data collected that comprises four datasets. The plants, time period, number of images, location and extra data are listed for each dataset.

Stress	Plants	Time	Images	Location	Dates	Multi view data
Spider mite	10 Elsanta, 10 Fenella	24 hours imaged daily	40 (20 x 2 days) 132 bands	NIAB EMR glasshouse	15th-16th June	-
Mildew	13 Elsanta, 12 Fenella	16 days -imaged two day intervals	225 (25 x 9 days) 132 bands	NIAB EMR Cabinets	22nd June - 8th July	-
Drought(E)	21 Fenella	5 days - imaged daily	105 (21 x 5 days) 132 bands	NIAB EMR glasshouse	13th - 23rd July (imaged 19th-23rd)	-
Drought(N)	6 Sonata, 12 BSP14, 8 Malling Centenary, 12 Hawaii4	6 days, 6 days, 5 days, 6 days	123, 812 bands	UoN Sutton Bonington campus	13th - 25th August	Pair + 3 - 10 images per plant

Observations from the datasets

The spider mite data is not being considered in this project. The experiment had to be shortened because the plants were infected with too many spider mites (+50) which resulted in the plants displaying visible signs of stress just 24hrs later. The data includes the control plants and infected plants before infection and at 24hrs. The aim was to collect more data before visible signs appeared.

Also, the data collected at EMR contains what appears to be a camera misalignment problem. It is not exactly known what the problem is but from contacting the manufacturer they believe it to be a misalignment problem with the sensors which affects the hyperspectral data and causes a shift noticeable when imaging a 3D object.

As the plant moves along the y axis the camera captures a line in the z axis, however there is an apparent 3D shift as the plant is moved appearing as though the plant has moved in each frame of the hyperspectral wavelengths.

To summarise:

- Spider mite - Biological problem: The over saturation of spider mites during the infection process resulted in visible signs appearing at the next imaging time point.
- Mildew and Drought(E) - Technical problem with the HS camera: The images included a visual shift in the wavelengths. This created a 3D effect in the images due to the misalignment in the camera sensors. The spectral signatures were affected.

While the first three datasets were captured at NIAB EMR, the Drought(N) dataset was collected at Nottingham with a different camera. Drought(N) also includes stereo pairs of hyperspectral data and Multi view colour images from the side of the plants.

Table 3.4: Table of the dataset selected after consideration of the data

	Plants	Time	Images	Location	multi view data
Drought(N)	6 Sonata, 12 BSP14, 8 Malling Centenary, 12 Hawaii4	6 days, 6 days, 5 days, 6 days	123, 812 bands	UoN Sutton Bonington campus	Pair + 3 - 10 images per plant

Image collection

The strawberry varieties for Drought(N) are Sonata, BSP14, Malling Centenary and Hawaii4. There are some images of each variety in Figure 3.7. This displays two plants of each variety with the aim to show the appearance and the variety within and between varieties. For example, Figure 3.7e-f are Hawaii plants and it can be seen that the size of the plants are different. Most of the Hawaii plants that were imaged (10) are large and complex with a few small sized Hawaii plants (2). BSP (a-b) and Hawaii (e - f) have small sized leaves but BSP leaves have a darker shade of green. Sonata and Malling Centenary are generally medium/large in size. The exception is (d) which is a Sonata plant. The Sonata plants imaged generally had flat leaves whereas Malling Centenary had lots of leaves that were curled and overlapping.

To summarise:

- The BSP14 plant is generally small which might retain water and last longer during water stress. There was a drought induced study that included BSP14 (*Fragaria chiloensis*) and the results stated that there was no difference between the control and drought plants for the first four days [88].
- Sonata cultivar was created from Elsanta and Polka. Sonata plants are medium/ large in size and the leaves are large (in comparison with BSP14). The leaves overlap and the plant is quite dense. However the leaves during the well watered period appear flat and they face the camera.
- Hawaii-4 is known as a wild strawberry plant. These are not a commercial fruit producing plant, the strawberries produced are small and white/yellow in colour and the leaves are very small.
- Malling Centenary was created at NIAB EMR in 2013 to compete with Elsanta. This plant is similar to Sonata in size, the leaves are dense and the the plants grow fast.

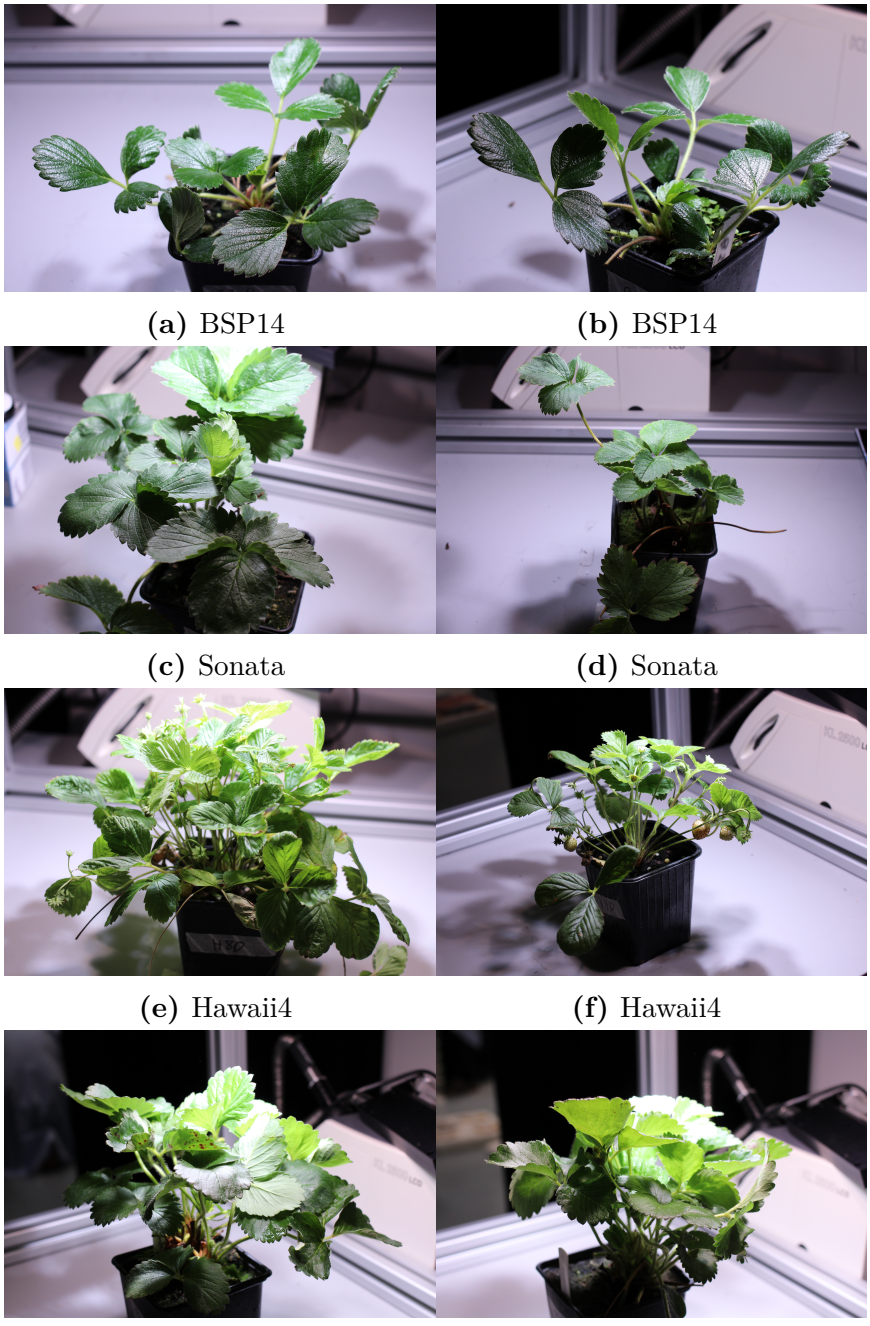

(a) BSP14 **(b)** BSP14

(c) Sonata **(d)** Sonata

(e) Hawaii4 **(f)** Hawaii4

(g) Malling Centenary **(h)** Malling Centenary

Figure 3.7: Images of the four strawberry plant varieties to display the appearance and architecture of the plants. BSP14 and Hawaii have small leaves although BSP14 remained quite small whereas Hawaii have a variety of sizes (e - f). Sonata have large leaves and are a medium size. Malling Centenary are of a similar size but the leaves are curling (g -h)

Selection of varieties

The BSP14 plants were selected due to the small organised plant structure making

imaging easier. This is a good control variety due to evidenced early drought tolerance. BSP14 would also be good for computer vision techniques development because the leaves are small, the plant in general is small and the leaves are spread out.

Sonata plants were also selected because they originate from Elsanta, and the leaves are relatively flat which is helpful for hyperspectral data.

Hawaii-4 and Malling Centenary have not been selected for further study. Hawaii-4 is similar to BSP14 in leaf size however the variety is a wild strawberry plant. Most of the plants imaged are like Figure 3.7e, dense compact leaves. This is not suitable for sections of the later work, and with the amount of leaves overlapping any ground truth labelling would also be challenge. Malling Centenary is a good option, it is similar to Sonata; however the leaves have a slight curve, overlap more and look compact as well as large.

From this point onward drought(N) captured in Nottingham will just be known as the drought dataset.

3.1.6 Data calibration

Calibrating the images

The hyperspectral data needs to be calibrated to ensure the images produced are adjusted for the lighting; hyperpspectral camera software may have this option, but if it does not then the data can be calibrated after it is captured. The lighting is calibrated using a known white balance target, which is captured by the camera system in the same imaging period as the plants. An example image of a small white balance with a BSP14 plant is included (see Figure 3.8), however a larger white balance board was imaged at a similar height as the top of the plants. This target will reflect a known percentage of light over the complete spectrum, for example 99% across the entire working spectrum of the camera. The non-uniformity of illumination can be corrected for by dividing the observed data by the captured white balance data (see Figure 3.9). Additionally, the system must be corrected for electrical noise present from the sensor in the absence of light (called dark current). This is usually carried out by taking an image with the camera in the absence of any light, and using the resultant low-level noise readings to adjust future measures. [89]

$$X = \frac{S - B}{W - B} \tag{3.1}$$

Where X is the resulting calibrated image, S is the input image, B is the dark current image and W is the white reference. [90]

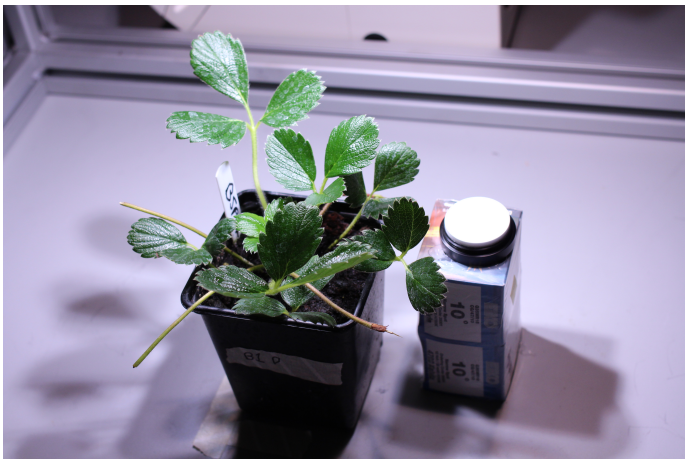

Figure 3.8: A BSP14 plant and small white balance disc. This is an example, a larger white balance board was placed at a similar height to the top of the plant and imaged for the calibration process.

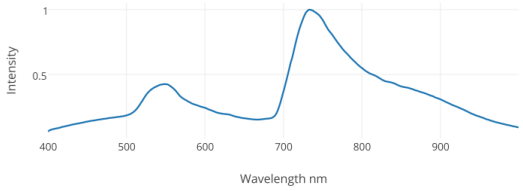

(a) Uncalibrated signature of a strawberry plant

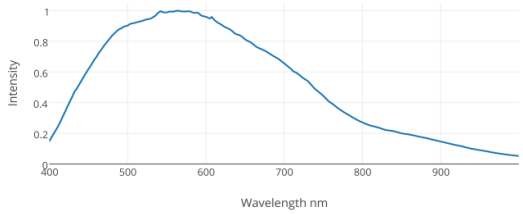

(b) The white balance signature

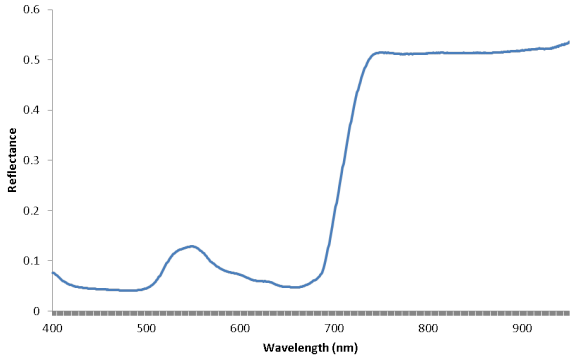

(c) Calibrated signature of a strawberry plant

Figure 3.9: (a) The signature of the strawberry plant before the image is calibrated; (c) The signature after it has been calibrated using the light colour spectrum (b). (c) The small peak at 540nm represents green, the sharp incline is called the red edge and the larger peak is the near infra-red range.

The data collected at EMR NIAB already had this process included in the camera software. Therefore the white balance needed to be imaged and the dark reference involved turning the lights off and covering the lens to capture a dark current image. Then these two references were included in the calibration formula. However the data collected from Nottingham did not include the white balance calibration. In this case the white balance reference is divided by the images to calibrate the data post capture.

3.2 Part 2: Initial data inspection and manual analysis

3.2.1 Manual selection of data

To provide a ground truth analysis, a manual approach to data extraction was explored. Figures 3.10a-b were produced from the Sonata plant images in the Drought-2 experiment to see if there were any noticeable differences between the drought and control plants. Six Sonata plants were imaged, three were control plants and three were drought-induced plants. The image analysis software package 'FIJI' [91] was used to manually select three leaves from the plants and take the average per day over three days for the control and drought plants. The control and drought graphs are displayed in Figure 3.10, the spectral signature have been smoothed automatically in 'FIJI ', the hyperspectral data unfiltered would be noisy in the near infrared range of the spectrum and this can be seen in Chapter 6.

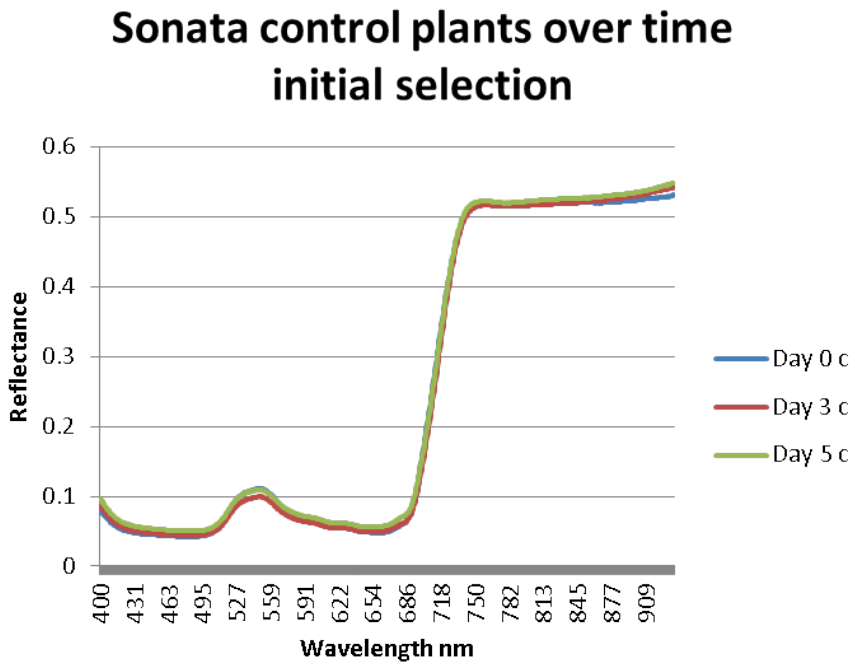

(a) Control Sonata plants spectral signatures over time.

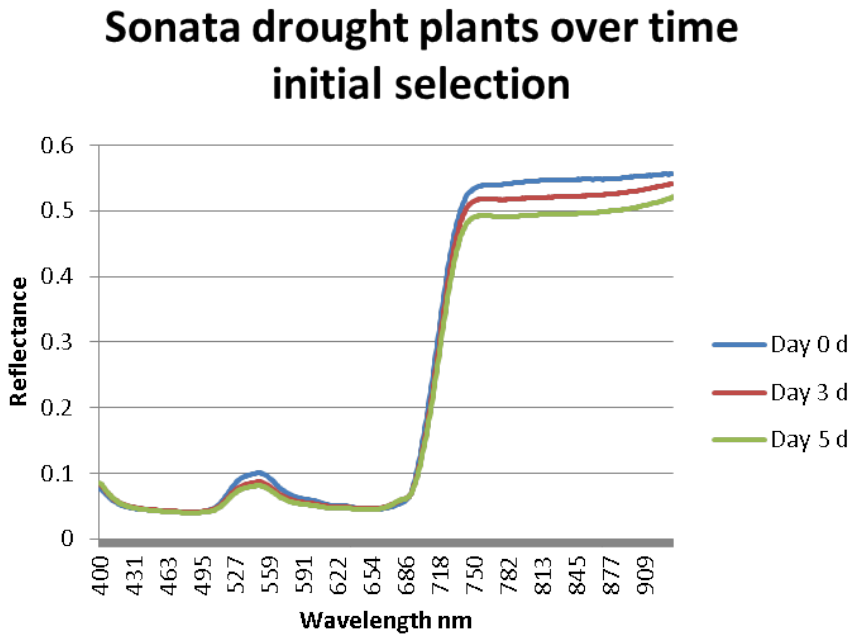

(b) Drought sonata plants spectral signatures over time.

Figure 3.10: The manual selection of control and drought plants spectral signatures plotted over time to display the difference between the two datasets. Sonata was used to see if a difference could be seen, as BSP14 has a suspected tolerance to drought.

From Figure 3.10 there are minimal differences between the reflectance of the control plants and some difference for the drought plants but it is hard to identify therefore more complex analysis is required. This data is manually selected from the plants, the next step would be to try and automate the process of selecting certain leaves and if it is possible to get better quality measures.

3.2.2 First automated analysis using the hyperspectral data alone

The hyperspectral data is both a spatial image and large complex data reflectance set. The first approach tried was to check what would happen if the hyperspectral reflectance data was analysed in isolation of spatial image information.

Using a mathematical approach to separate the images into endmember materials to select leaves

Refering back to the literature review in Chapter 2, Non Negative Matrix Factorisation (NMF) has been explored as a technique that only uses the spectral information from the data without spatial information. NMF is a matrix factorisation approach that minimises the error between the input and two expected output matrices. For hyperspectral data one output matrix is the expected hyperspectral signatures for N number of endmembers and the second matrix is the probability of the pixel being a part of that endmember (abundance).

There are three colourmaps produced from the matrices of the results (See Figure 3.11). The positive point is the technique is autonomous and requires no inputs after the initial parameters are set. The negative points are the number of endmembers need to be selected and also the data can be quite noisy. This means the results would be very hard to interpret between diseased and healthy for the data collected. But the technique is good at separating the endmembers/different types of material from the signatures. This separation can then be used to examine spectral profiles for underlying materials. A labelled figure with colour sections of the plant has been included for clarity of the plant materials (Figure 3.12).

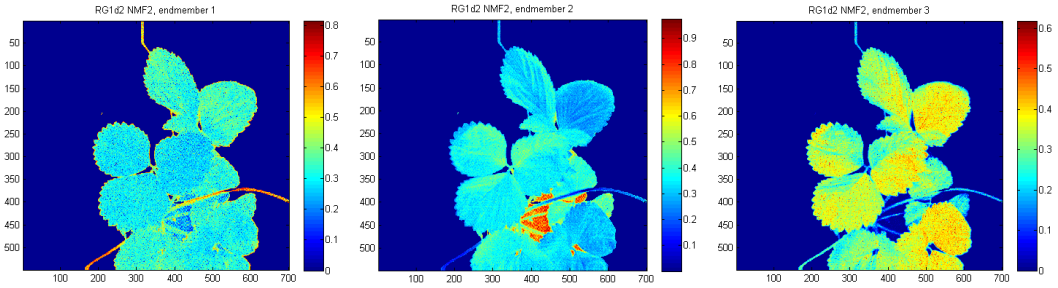

(a) NMF 1 of three endmembers - stolon **(b)** NMF 2 **(c)** NMF 3

Figure 3.11: The colourmaps of the three endmembers/parts of the plant seperated. Red is the endmember and blue is not the endmember. Figures where (a) is the first endmember the stolon, (b) is the second endmember the soil and (c) is the third endmember the leaves.

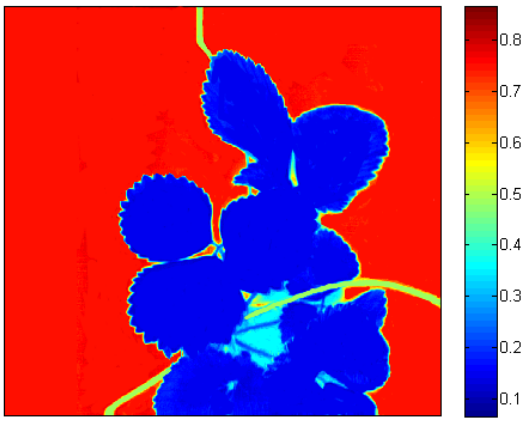

Figure 3.12: Labelled figure of the plant; dark blue is the leaf, light blue is the soil, green is the stolon and red is the background.

Two graphs are displayed for comparison between the ground truth spectral signatures (selected manually using FIJI) and the NMF produced spectral signatures (Fig 3.13). Figure 3.13b is the output from NMF; clearly it is very noisy. Near infrared wavelengths in the spectral data are very noisy to start with and no post processing has been applied.

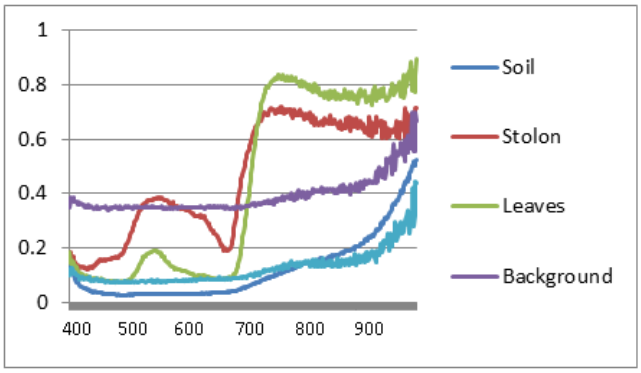

(a) The ground truth spectral signatures of the endmembers (materials)

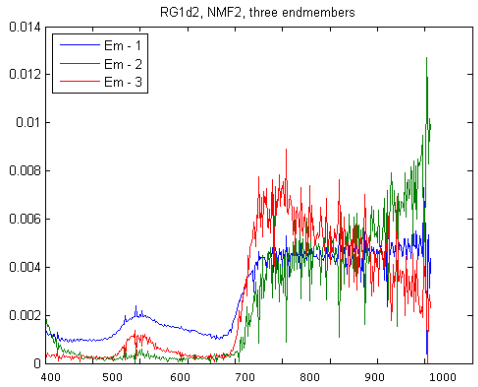

(b) The spectral signatures produced from NMF.

Figure 3.13: The two graphs display the spectral signatures of the endmembers (materials) in the image. Figures where (a) is the ground truth spectral signatures and (b) is the graph produced from the technique. Endmember1 is the stolon, endmember 2 is the soil and endmember 3 is the leaves.

The results from the technique are not good enough to analyse for small changes in the spectrum, while the flowers, leaves, stolon can be detected these have different shape spectral signatures whereas individual leaves have almost identical signatures. So are all clustered into one class. Therefore identifying individual leaves is not possible. Also the hyperspectral data is quite noisy and this technique will not work that well with the subtle changes that will occur in the spectral signatures as the plants become stressed. Therefore spatial information and computer vision techniques are going to be explored.

3.2.3 Conclusions from the initial tests

The Sonata plants were imaged over several days until visible signs of drought appeared, the hyperspectral data is displayed in Figure 3.10. From inspection of Figure 3.10 the control plants reflectance responses are very similar and the drought induced plants display signs of reducing in the green and near infra red wavelengths over the experimental period. The plants were imaged under the same conditions and camera settings. There is a suggested difference in reflectance over time, however it is hard to draw conclusions from this data alone.

The NMF analysis involved automatically selecting all the plant data and separating it into the different materials. When the number of materials increased the data was split into sub-categories of the same materials which produced noisy data. Also stressed areas would be difficult if not impossible to determine using this method alone.

NMF uses the hyperspectral data alone; this is a good test to see what could be detected from an automated process that uses just the colour (no spatial) information. Segmenting leaves based solely on colour when the colour is expected to change as part of the stress is not ideal. Geometric, vision based approaches are therefore required for this dataset. Developing such approaches is the foundation of Chapters 4 - 6 in the rest of the thesis.

In order to develop such methods, labelled ground truths are required which will be detailed below.

3.3 Creating 2D ground truth data

Before any methods can be evaluated ground truth data is needed. In this case leaf segmentation is needed. The process of creating the ground truth data manually involves drawing around the edges of leaves. This will create two sections, inside and outside the leaf. A mask image is used to label regions. The area inside the leaf is set to 1 and outside the leaf is set to 0 to create a binary mask. However plants usually have several leaves, therefore the leaves are labelled by a numerical value; background regions are labelled 0 (black) and leaves are labelled with ascending numerical values.

3.3.1 Hyperspectral images

The hyperspectral plant data is calibrated with the corresponding white balance data to normalise the illumination. Then a slice is selected, 700nm, to display the leaves. The leaf is manually selected, this would save the binary mask as 0 and 1, then the process is repeated for each leaf. Each leaf that is selected is then added to the mask matrix until all the leaves have been selected. The resulting image can be viewed as a colourmap where the leaves are displayed in different colours (index values) or they can be displayed with the boundary of the leaf on top of the 700nm image. Figure 3.14 shows four different views of the plant data from the initial grayscale image (a) to the ground truth selected (c) and the corresponding colourmap (d) to visually see the leaves. The colourmap index image and a binary image are saved to be used for the evaluation of the images in Chapter 4.

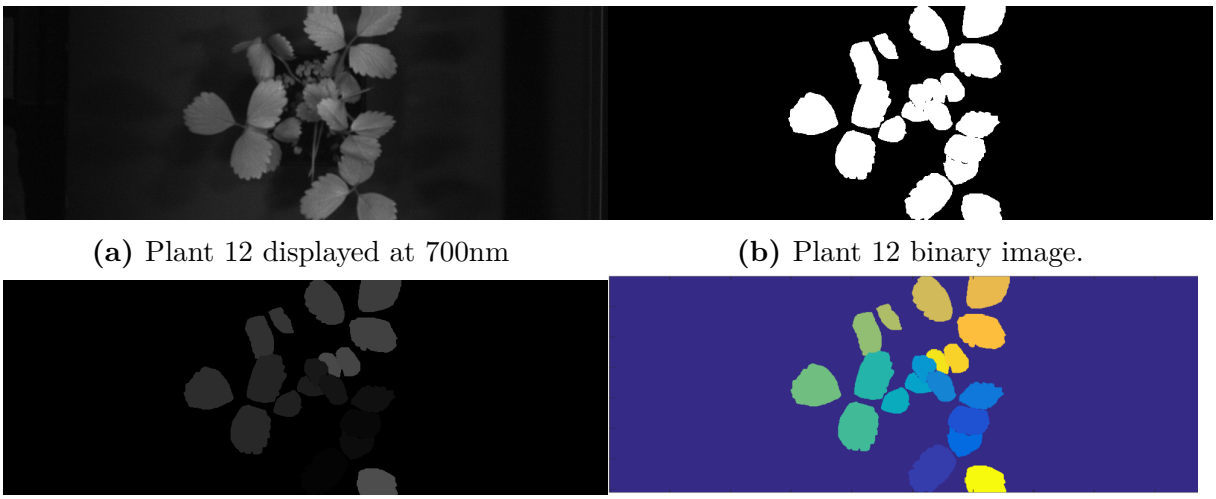

(a) Plant 12 displayed at 700nm

(b) Plant 12 binary image.

(c) Plant 12 ground truth, each leaf has a number id. Displayed in grayscale.

(d) Plant 12 colourmap.

Figure 3.14: Plant 12 - the ground truth data.(a) is a grayscale image of the plant from the hyperspectral data, used to draw around the leaves. (b) is a binary image. (c) is an example of the ground truth data that is used, it is dark because the leaves are numbered from 1-N. (d) shows a colourmap applied to (c) to help visualise the leaves.

The datasets with 2D ground truth that have been created are all Sonata and BSP14 plants for days 0, 3, 4 and 5. Malling Centenary also has 2D ground truth images however this plant variety is not used.

3.4 Conclusions

This Chapter covered the data capture, data selection and initial analysis. Several plant stress imaging protocols were developed and used. Three stress datasets were captured and compared with regard to suitability for analysis. Drought is promising to investigate because the stress is abiotic and it is consistent with how the typical hyperspectral signature profile should appear for healthy and stressed plants.

Initial visualisation shows some difference in reflectance over a drought experiment but there are only small changes visible. Colour-based clustering was explored as a way to identify leaf data, but was not found to be reliable. More sophisticated methods are needed to identify and select leaves for data analysis, due to the challenges of hyperspectral image data. These will start with a segmentation using a leaf shape model, explored in Chapter 4.

Chapter 4

A model based approach for segmenting leaves in hyperspectral images

This Chapter will describe the development of computer vision techniques in order to segment individual leaves. Chapter 3 initially analysed selected leaves using manual analysis which indicated a slight difference in the spectral signatures for the drought plants compared to the control plants however the leaf positions, shadows, overlapping and depth can all affect the spectral signatures. The logical next step is to find the individual leaves to be able to analyse them separately. Additional analysis would be needed later in order to determine which of these segmented leaves should be selected to take hyperspectral measures from.

The image data is quite complex, with texture, occlusion of some leaves and lighting changes. The main aim is to segment individual leaves in these challenging images automatically with no input from the user.

Separating the whole plant from the background is relatively easy for this data (when the background is suitably different) by using colour to seperate the plant from everything else - k-means, gaussian mixture models or the hyperspectral profiles can be used. NMF in Chapter 3 could separate the plant from the background. It becomes more challenging when the individual leaves need to be segmented, as described in Chapter 2.

The hyperspectral images collected here have distinctive strawberry leaf edges and the main vein of the leaves are mostly visible. After the hyperspectral images

were captured a second set of SLR images were also captured both to provide high resolution option and for a reference. Hyperspectral data is relatively low resolution compared to SLR, and also tends to have a narrower focal plane which is even more of a challenge. The hyperspectral dataset will be used for the analysis but both datasets will be displayed to explain the proposed pipeline (see Fig. 4.2).

4.1 Leaf model pipeline

The leaf model is an automated pipeline to find the leaf shape segmentation for the leaves in the images. This process looks at finding the leaf centre points and the veins in order to place a leaf shape template over the leaf and rotate it in the same direction as the leaf. The shape template is then incorporated into the level set to find the final leaf segmentation. Figure 4.1 is a flowchart of the pipeline. The numbers in the flowchart correspond to the algorithm.

The algorithmic steps are:

1. Find the individual leaf centre point,
2. Find the veins
3. Create a mean shape template
4. Place the shape template in the local area of the leaf using the leaf centre points (from step 1)
5. Rotate the shape template using the veins (from step 2)
6. Apply the level set with the shape incorporated. (from step 4)

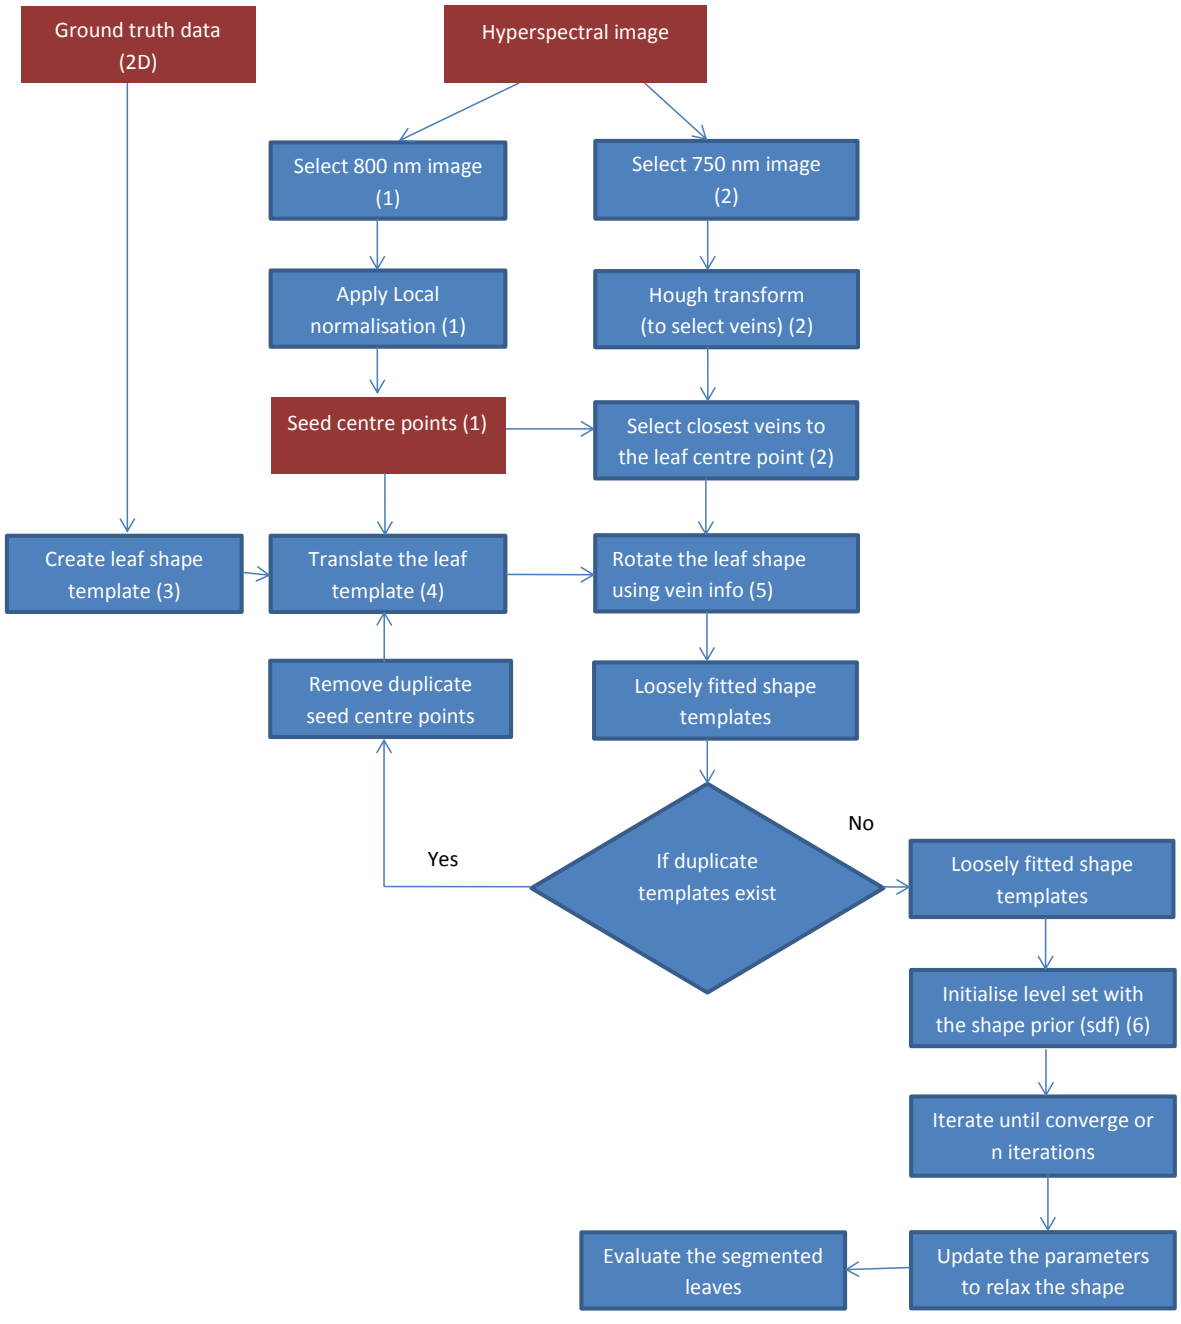

Figure 4.1: Flowchart for the hyperspectral leaf segmentation pipeline presented here in Chapter 4. The numbers correspond to the algorithm steps.

The dataset being used for the model is the hyperspectral images collected in Nottingham for the drought dataset; specifically BSP14 and Sonata.

(a) RGB data from DSLR

(b) Hyperspectral data displayed in RGB

Figure 4.2: Figures where 4.2a is DSLR images and 4.2b is the hyperspectral images displayed in RGB. The hyperspectral images are of relatively poor resolution and quality (b)

4.1.1 Finding the leaf centre point using local normalisation

The first step in the leaf model is to find the leaves centre point. There is a potential method called the distance transform which measures each pixels distance to the boundary of a binary image. This however tends not to work that well when there are multiple leaves together and overlapping.

Local normalisation takes a small window of the image and applies the local mean and variation to the pixels then the window is moved and applied as a filter to the whole image. By applying a Gaussian smoothing kernel first and then local normalisation the image is broken into regions. Finally a size threshold is applied to remove small groups of noisy pixels and keep the large clusters which represent the leaves. In images with severe overlapping leaves duplicate seeds will be in a leaf which can be removed later.

Images of three strawberry plants with different amounts of leaves have been used to display the output from local normalisation (see Figure 4.3). (a) shows lots of leaves but the leaves are bright and only slight overlapping. The centres of the leaves appear to work well apart from two in the background that are partially occluded. The middle image (b) is slightly darker in intensity than the top image and the leaves are closer together. Here local normalisation has not worked as well and the seed centre points are closer to the edges of the leaves. The bottom image (c) has more leaves and a few leaves are out of focus. Multiple seed centre points have been placed on three leaves while the other seed centre points are near the edges of the leaves.

The hyperspectral images are noisy especially in the near-infrared section, and the RGB images are higher resolution. However the near-infrared section of the hyperspectral images are bright because the plants reflect a lot of light for this wavelength range which helps the local normalisation. Using NIR is not possible on RGB images, but optimising the process to use NIR(eg. 800nm) data where available works well.

Finding the leaf centre points (ie. seed locations for the rest of the leaf segmentation algorithm) is evaluated further in 4.3.1.

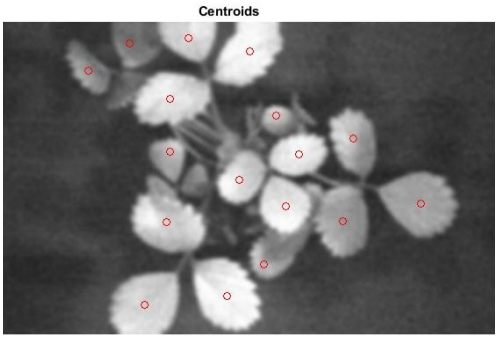

(a) Hyperspectral image displayed at 800nm. The circles are the automated seed centre points. There is a good contrast between the leaves and background.

(b) RGB displayed in grayscale with the automated leaf centre points. Although for this image they have not found the centre, they are near the edge of the leaves.

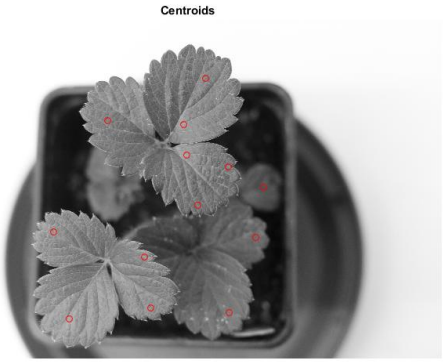

(c) RGB displayed in grayscale with the automated seed centre points. The seed centre points are not in the centre.

Figure 4.3: Greyscale images of strawberry plants with the centre points displayed on the leaves. (a) leaves are spread apart and bright (800nm), (b) leaves are overlapping but there are only a few leaves, (c) multiple leaves, close together. The automated seed centre point detection works well with the hyperspectral data however they appear to not work as well with the RGB high resolution images.

4.1.2 Finding the veins on the leaves for shape template orientation

After finding the centres of the leaves, the next step is to find the veins. The veins will help to rotate a shape template to the same direction as the leaf orientation. The veins in strawberry leaves have a distinctive architecture, there is the main vein that is direct from the stem to near the end of the leaf, then the secondary veins branch out almost symmetrically at about 30 degrees each side of the main vein. (see Figs. 4.3 and 4.4)

A way of finding the veins is by using the Hough transform [92]. This method can find straight lines or even curves in the image. First an edge detector is applied to the image, this will find lots of lines on the image. Then points are mapped from the image to curves in the Hough space. The peaks of the curves represent parameters of potential straight lines. There are parameters to set in the Hough transform to control the length of the lines detected, and the number of lines. Figure 4.4 displays the lines on a strawberry plant image found using this approach.

Figure 4.4: Lines placed on the RGB image from the Hough transform. The image is blury in the background therefore the leaves in the distance are out of focus and not as many veins are found. Found lines are in green; line end points are indicated by the red/yellow Xs.

Many lines are found. It is necessary to reduce the amount. Lines close to the leaf centre points are preferable and allow for the general leaf direction.

Figure 4.5 displays two automated leaf centre points with the three nearest lines.

The leaf on the left works well however because the leaf on the right is out of focus no lines are found on the leaf and instead the lines are from nearby leaves. For this project it is ok that the leaf would be disgarded because it is out of focus, in the background and slightly occluded, so unsuitable for measurement purposes anyway. This is to display problems with the current method.

(a) Veins for the bottom left leaf

(b) Veins for the lower, out of focus leaf

Figure 4.5: Lines closest to the centre points are selected and displayed. (a) The selected leaf is the bottom left with the three closest lines selected successfully. (b) The selected leaf is on the right in the background (white circle), which is out of focus; no veins are detected successfully for that leaf.

4.1.3 A shape template for leaf segmentation

The automated leaf centre points and the veins will both contribute to placing a leaf shape over the leaf. However a leaf shape template needs to be created first. Strawberry plants can vary in size and shape, two varieties that will be used throughout the thesis are BSP14 and Sonata. BSP14 leaves have a small circular shape while Sonata leaves are large, irregular and they have a pronounced jagged edge. One way is to use a particular ground truth strawberry leaf as the leaf shape. However the shape may be too big or too small and the shape could certainly be considerably different to a second variety.

The next idea was to use a mean leaf shape; this approach takes several ground truth leaves, combines the binary masks and calculates the mean shape.

One mean shape template would not account for the large shape variation of leaves between the varieties. For this dataset, BSP14 and Sonata, there are two mean shape templates to account for the large difference in the leaves. The aim is not

to match the shape template perfectly but to have a general leaf template where the leaf is.This template will form the basis of the level set initialisation later on.

Figure 4.6 shows the rotation and translation of two leaves. The shape template is rotated using the vein information located by the Hough transform; the mean line direction, then the shape template is translated to the automated seed centre points from local normalisation.

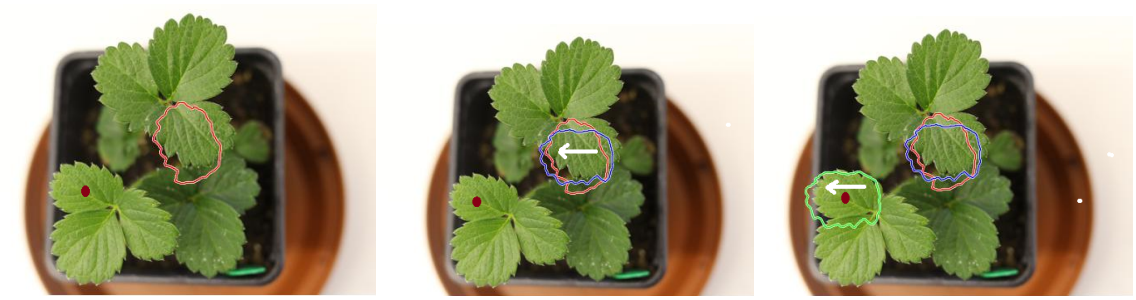

(a) Initial shape (red) **(b)** Rotation (blue) **(c)** Translation to seed point of desired leaf (green)

(d) Initial **(e)** Rotation **(f)** Translation

Figure 4.6: Two examples displaying the shape template being rotated and translated using the seed placement and vein structure described above. Left: The red shape template starts in the middle of the image. Middle: The blue shape template has been rotated from the red template to the vein orientation of the leaf. Right: The green shape template has been translated to the leaf. Dot on the target leaf centre to show where the shape template is being moved to. Arrow to show the direction the template has been rotated to.

4.1.4 Leaf model up to the shape template placement output: example for one complete image segmentation

At this point the leaf centre point and the leaf veins have been located for the majority of the leaves. If the leaf is dark or blurry then it might not have been selected. Then the mean shape has been created, rotated and translated to where the leaf should be.

Figure 4.7 displays the outputs for each stage of the leaf model. First the leaf centre points are detected, then veins on the leaves are located. Both of these steps are used to place the shape templates over the leaves. The shape templates are placed over the leaves (Fig. 4.7c), there are overlapping templates which means there are duplicate seed points.

Template accuracy forms part of the final boundary evaluation in Section 4.3.3.

(a) The automated leaf centre seed points detected (green circles).

(b) Veins located on the leaves using the Hough transform (green lines, crosses represent line termination points)

(c) The placed shape templates after they have been translated to the leaf using the centre points and rotated using the veins.

Figure 4.7: The output from the process of the model. (a) Output from the automated leaf seed centre points, the green circles represent the points (b) the lines/veins are found using the Hough transform and finally (c) the shape templates are placed using these steps to rotate and translate the mean shape template. The placed shape templates in (c) are used in the level sets to find the final leaf segmentations.

The next step is to incorporate the shape template into a level set approach to segment the individual leaf boundaries.

4.1.5 Level sets and the shape prior

The Chanvese model, introduced in the literature review (Chapter 2.1.1) is used because it is good for finding weak edges which is needed when the leaves overlap and the difference in the pixel intensity is small [27].

One problem associated with the use of level sets, is that it will find any connecting regions of the image instead of individual leaves. To prevent this there are different versions of level sets which constrain the contour to a specific area in order to find the nearest edges/contours. However, the leaf boundary is hard to see with the human eye therefore additional steps are needed. When the image information alone cannot find the leaf, then incorporating prior shape information is a possible solution. This involves creating a shape model that resembles a strawberry leaf and incorporating it into the level set. The template created in Section 4.1.3 can be used as a shape prior and incorporated as so:

$$F(c1, c2, \phi, \phi o) = F(c1, c2, \phi) + \alpha shape(\phi, \phi o) \tag{4.1}$$

$F(c1, c2, \phi)$ is Equation 2.10 in Chapter 2.1.1.

$$shape(\phi) = \int (\phi - \phi o)^2 dx \tag{4.2}$$

Where ϕo is the shape model, a signed distance function. Equations (4.1) and (4.2) represent the updated level set to include the shape.

This shape incorporated level set model was developed by Cremers, Sochen and Schnorr [93]. They use the Mumford Shah model and embed the shape into a signed distance function. There are two options to incorporate the shape into the model. The first option is to include the shape into the evolution equation and the second option is to add the shape energy into the signed distance function.

Here, the shape is included into the functional. Also the shape template is used with the Chanvese model rather than the Mumford Shah because the leaves overlap and the level set needs to be sensitive to weaker edges.

An additional inclusion into the level set approach is the reinitialisation of the signed distance function to update the model, which allows for customisable parameters. The customisable parameters would allow the objective to be flexible and adjusted over time and after iterations with the shape being the primary objective then the parameters could weaken the shape force and strengthen the edge forces. This could improve the segmentation of edges rather than being smoothed from the pull of the sharper shape forces.

4.2 Evaluation and Results

Two stages of the model are evaluated: the leaf centre point placement accuracy and the final leaf shape segmentation.

4.2.1 Seed centre point location accuracy

The automated seed centre points are used to initially place the leaf shape template in the same place as the leaf, this is then incorporated into the level set to find the leaf shape segmentation. The centre points will be measured against ground truth centre points to determine the accuracy of the automated leaf centre points.

To evaluate the seed point centres a ground truth is needed. The ground truth was the leaf masks described in Chapter 3, where each pixel representing a leaf on the plant is labelled from 1 to n. The centre points have been selected by finding the centre of each binary leaf mask - the average point of the x and y points in the leaf area. Figure 4.8 displays the ground truth centre point in Plant 12 (BSP14, plant 12, day 3).

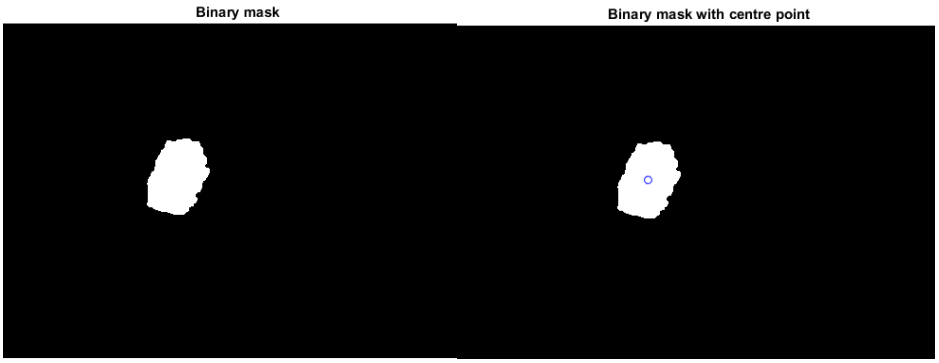

(a) Binary mask of a leaf from ground truth (Chapter 3)

(b) Binary mask of a leaf with the ground truth centre point

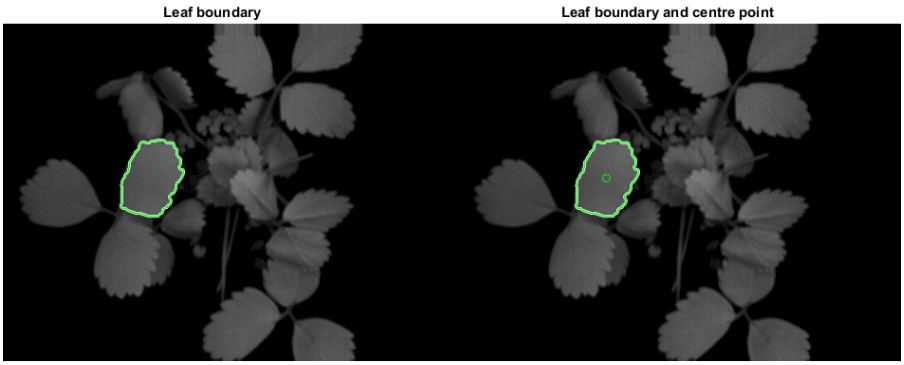

(c) Boundary from (a) overlaid on image

(d) Leaf boundary and ground truth centre point overlaid on image

Figure 4.8: BSP14 Plant 12 ground truth process for selecting the leaf centre points.

An example of the ground truth centre points and the automated centre point locations are compared in Figure 4.9. Panel (c) displays the difference in location by plotting both of the centre points on the same image where the yellow circles are the ground truth centre points and the green circles are the automated centre points. The accuracy is measured by using the Euclidean distance to find the distance between two points. The results display the average plant accuracy in pixel units across all leaves as well as the root mean squared error. The results are in Tables 4.3 and 4.5.

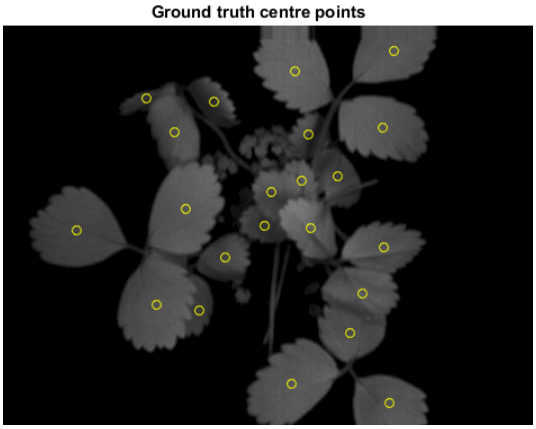

(a) Ground truth leaf centre points

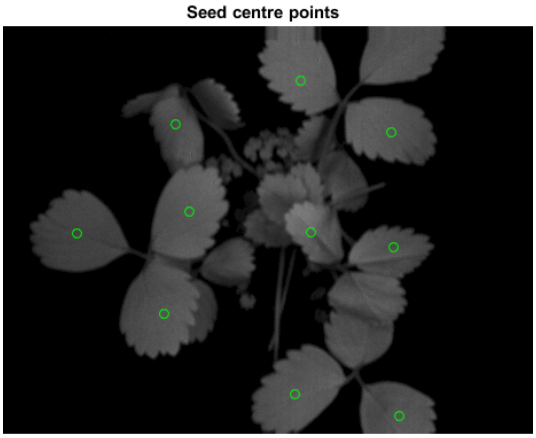

(b) Automatically detected centre points

(c) Comparison between the ground truth leaf centre points (yellow) and the automatically located centre points (green).

Figure 4.9: Example comparison between the ground truth centre points and the seed centre points for BSP14 plant 12. Some leaves have not been detected, this will be due to the brightness or blurriness of the leaves (out of focus leaves).

Table 4.1 displays placement results for 12 BSP14 plants. Accuracy based on

distance is calculated in 4.3.2.

- Plant - the plant number
- Number of leaves - the number of leaves that are on each plant - ground truth (GT)
- Number of seeds - the number of centre points automatically detected (auto detect)
- False positive seeds (FP seeds) - the automated centre points that missed the leaves completely
- Duplicate seeds - the number of automated centre points that appear on a leaf more than once.

Table 4.1: BSP plants seed information

Plant	No. of leaves (GT)	No. of seeds (auto detect)	FP seeds	Duplicate seeds
1	21	8	3	0
2	20	10	1	0
3	13	5	1	0
4	16	12	2	0
5	18	5	1	0
6	17	10	1	0
7	23	11	1	0
8	18	7	1	0
9	16	5	2	0
10	22	6	2	0
11	19	8	2	0
12	22	10	1	0

Sonata plants also need to be evaluated because the varieties are visually different strawberry plants and the overlapping crowded leaves means this data is more challenging. Note for the plant in Figure 4.10 there are duplicate seed points in a leaf. An example Sonata plant (Plant 13) is displayed in Figures 4.10a,4.10b and 4.10c. The leaf centre accuracy results are in Table 4.4.

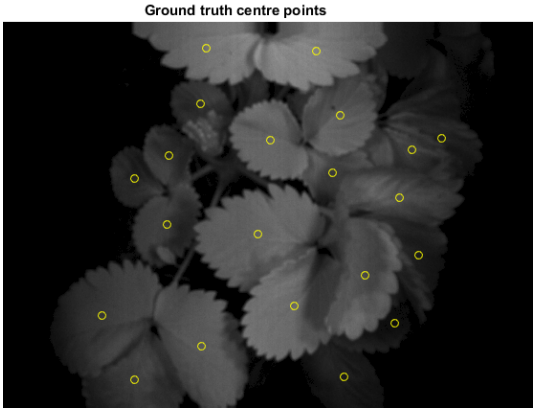

(a) Ground truth leaf centre points for a Sonata plant

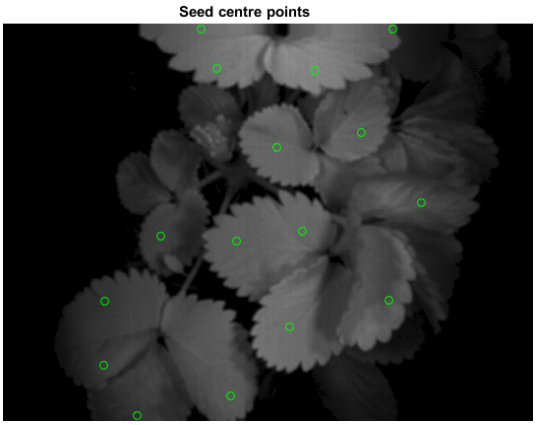

(b) Automatically detected centre points. There are some duplicate seeds (¿1 per leaf).

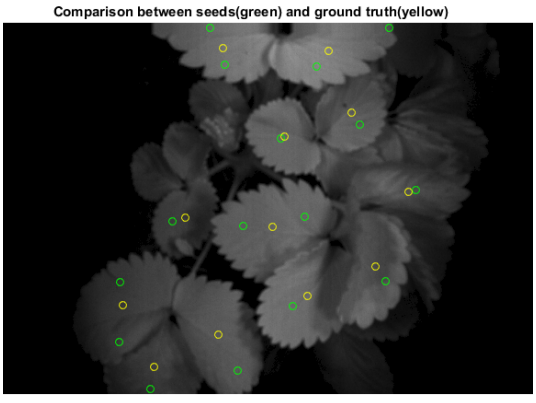

(c) Comparison between the ground truth leaf centre points (yellow) and the automatically located centre points (green).

Figure 4.10: Comparison between the ground truth centre points and the seed centre points for Sonata plant 13. Some leaves have not been detected, this will be due to the brightness of the leaves. As can be seen from (a), some leaves are very dark, even by eye.

Table 4.2 displays the same information as Table 4.1 for the Sonata plants. There are duplicate seed centre points which means some leaves have more than one seed per leaf.

Table 4.2: Sonata plants seed information

Plant	No. of leaves (GT)	No. of seeds (auto detect)	FP seeds	Duplicate seeds
13	21	12	0	4
14	16	6	1	1
15	21	10	0	1
16	15	10	0	1
17	21	12	0	1
18	17	9	0	0

The duplicate seed points have been removed if two shape templates are in the same area (by 80 percent or more). The original seed points were selected and combined by selecting the midpoint. Figures (4.11a to 4.11c) are Plant 13 again after duplicate seeds have been removed. The duplicate seeds are discussed in the results Section.

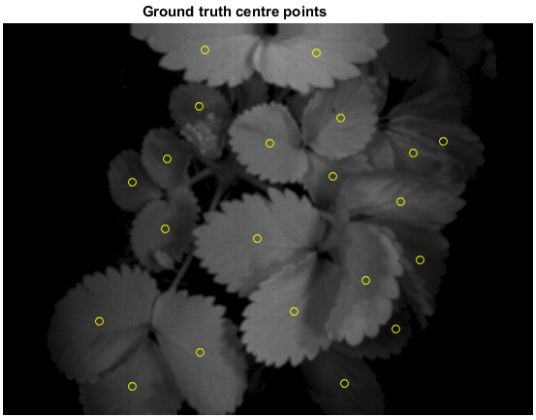

(a) Ground truth leaf centre points for a Sonata plant

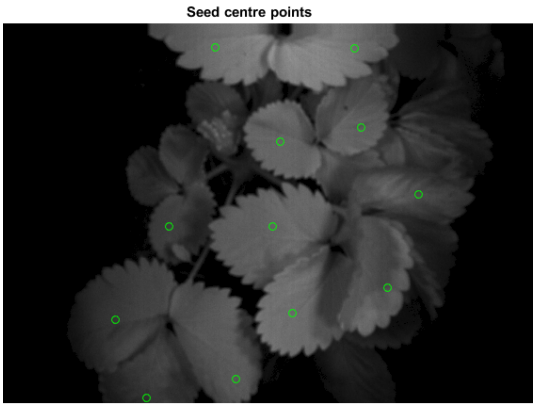

(b) Automatically detected centre points

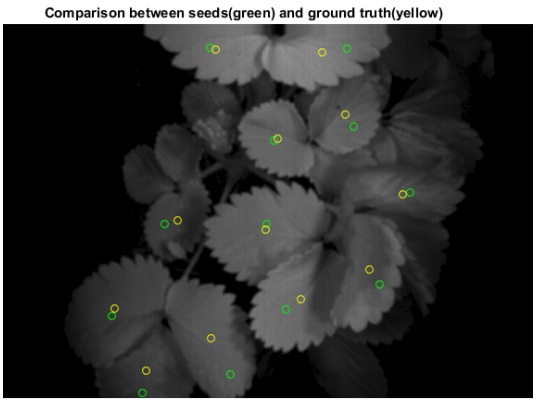

(c) Comparison between the ground truth leaf centre points (yellow) and the automatically located centre points (green).

Figure 4.11: Comparison between the ground truth centre points and the seed centre points. The duplicate seeds have been removed.

4.2.2 Automated leaf centre accuracy results

The average accuracy across all leaves on a plant for all of the automated leaf centre points are listed in Tables 4.3 and 4.5. This discounts missed seeds. There were no duplicate seeds for BSP14 however there were duplicate seeds for Sonata and this is in Table 4.4.

Table 4.3: BSP plants seed location evaluation

Plant	Average Distance (pixels)	RMSE
1	6.021	7.574
2	7.153	11.520
3	3.176	3.589
4	6.150	7.774
5	7.123	9.567
6	4.388	5.206
7	4.153	6.083
8	6.266	11.32
9	4.644	5.066
10	6.519	8.739
11	3.641	4.173
12	4.618	5.254

Table 4.4: Sonata plants seed location evaluation with the duplicate seeds still included.

Plant	Average Distance (pixels)	RMSE
13	21.327	30.506
14	26.019	31.248
15	27.465	31.246
16	18.712	22.086
17	23.835	27.779
18	15.506	19.045

The average plant accuracy for BSP14 is less than 8 pixels whereas the average plant accuracy with no duplicate seeds (from Table 4.2) for Sonata is 15-24 pixels from the ground truth centres. Three additional Tables in Appendix A lists example individual leaves for two plants of each variety to provide examples of the individual leaf accuracies, the ground truth and automated leaf centre points as well as the separate x and y deviation from the ground truth centre point. The Tables include the individual leaf information rather than the average plant results.

Table 4.5: Sonata plants seed location evaluation after duplicate seeds were removed.

Plant	Average Distance (pixels)	RMSE
13	16.548	19.903
14	19.146	21.779
15	23.519	25.095
16	16.769	19.296
17	20.386	23.210
18	15.506	19.045

The automated seed centre point detection works better for BSP14 than Sonata, this could be because BSP14 leaves are small, well lit and a similar shape. Sonata plant leaves are large, at various angles which changes the brightness and the leaf shape varies. The furthest leaf centre point error was Sonata plant 15 with a distance of 24 pixels (Table 4.5). An average Sonata leaf is 80 pixels in width and 96 pixels in length, therefore even the furthest centre point of 24 pixels (from the ground truth centre point) is still inside the leaf and therefore acceptable.

4.2.3 Level set final contour

The initial level set is placed using the leaf model pipeline. First the automated leaf centre points are located using local normalisation, and then the leaf template is translated to the centre point, then a few of the veins close to the centre point are used to rotate the leaf template. The leaf template is then incorporated as the initial contour for the level set, converted into a signed distance function.

Jaccard and Dice similarity scores are selected to evaluate the final segmented leaves. This uses two binary images and compares the similarity with a score between 0 and 1, where scores closer to 1 represents a closer match. If the value is 1 then both images match exactly, if the value is 0 there is no overlap between the two images; this is the lowest the value can be, representing the poorest fit.

Jaccard similarity score is the intersection over union:-

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \tag{4.3}$$

Dice similarity score:-

$$D(A,B) = 2 * \frac{|A \cap B|}{|A| + |B|} \tag{4.4}$$

Dice can also be written in relation to Jaccard (the same applies for the Jaccard index).

$$D(A,B) = \frac{2 * Jaccard(A,B)}{1 + Jaccard(A,B)} \tag{4.5}$$

Visual inspection

Figures 4.12 to 4.16 show the final leaf segmentation from the level set for several individual plants. Some leaves are missed from the final segmentation. This happens in the vein selection process if no lines are found near the centre point, also the centre point if the leaf is too dark. Referring to Table 4.1 in Section 4.3.1 the number of leaves that the ground truth finds are all of the leaves but the number of seeds found are a reduced number. Some leaves are missed from the automated leaf centre point detection if the leaves are very dark Eg. see Figures 4.9 and 4.10. This is not necessarily bad since the poorly lit leaves are usually near the bottom of the plant, or in a shadow. Therefore, those leaves would not be good for further analysis. The second place where leaves are removed are when the veins are not found for the leaf. Without the veins the leaf template would not be rotated correctly. This happens when the lines from the veins can not be detected due to noise or low resolution of the image. The leaves are rejected if no veins are located within the centred shape template.

The final resulting leaf segmentations generally segment the leaves well, however in some places where the leaves overlap and there is only a small difference in intensity between the leaves the segmentation misses the leaf boundary. Figure 4.13 displays a leaf where the segmentation has increased passed the leaf to locate a small section of the leaf below.

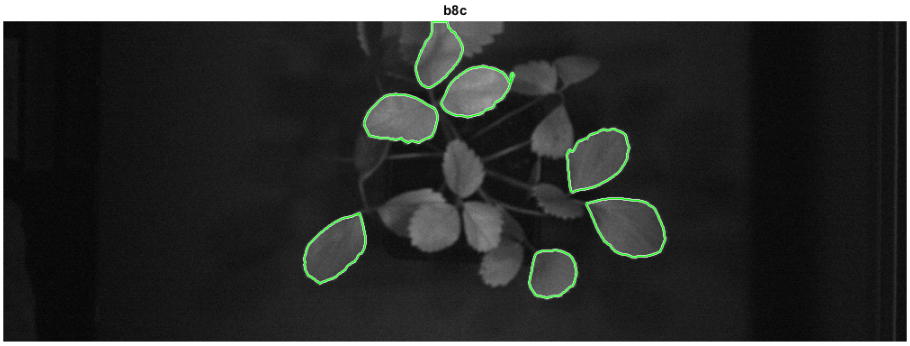

Figure 4.12: BSP14 plant 8 with the segmented leaves displayed in green

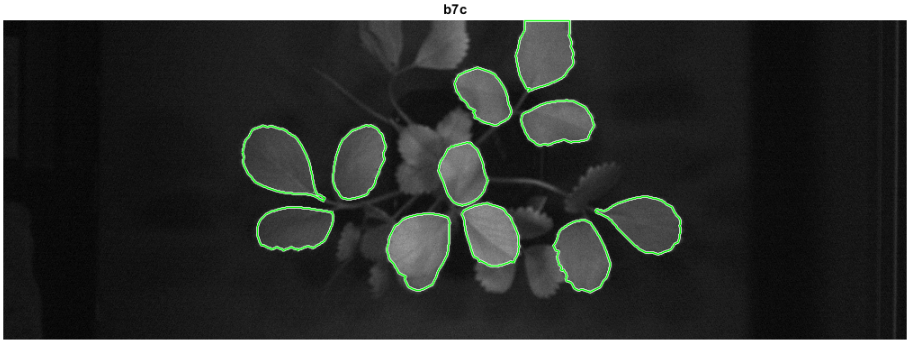

Figure 4.13: BSP14 plant 7 with the segmented leaves displayed in green. The bottom middle leaf (fourth leaf from the left) displays where the segmentation has over segmented because the difference in pixel intensity is small between the two leaves.

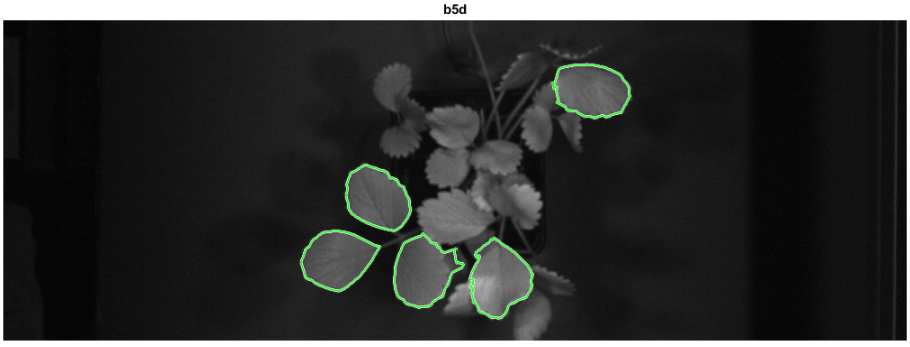

Figure 4.14: BSP14 plant 5 with the segmented leaves displayed in green

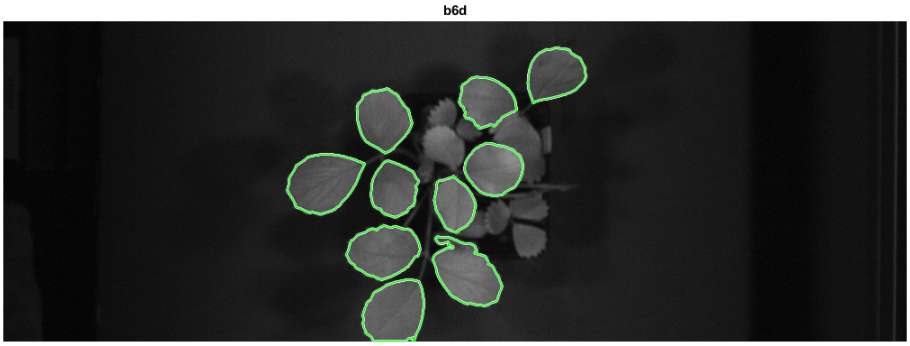

Figure 4.15: BSP14 plant 6 with the segmented leaves displayed in green

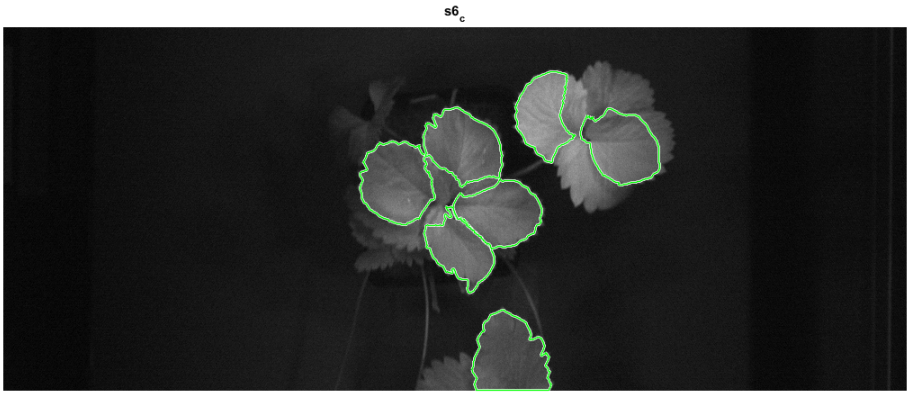

Figure 4.16: Sonata plant 6 with the segmented leaves displayed in green

Visual inspection summary

Poorly lit and occulded leaves are not detected by the pipeline at all. This is a positive result for this work because these leaves would not be good to analyse in Chapter 6.

The boundaries are a generally good fit, they will be measured in the next Section.

4.2.4 Leaf segmentation accuracy results

Table 4.6 displays the results for the leaf segmentation accuracy for boundaries which remain at the end of the pipeline. The plant overall accuracy average is the average of the metrics across all selected leaves on the plant. The segmentation accuracy for BSP14 is better than Sonata. The BSP14 leaves are small and have an easier shape. The leaves tend to be isolated without overlapping leaves, or

touching leaves. Sonata tend to be large leaves grouped together with overlapping and touching leaves and have a more complex shape. The averages from the Jaccard and Dice scores have also been included for an overall comparison.

Table 4.6: Leaf segmentation evaluation

Plant	Jaccard (0-1)	Dice (0-1)
4	0.862	0.922
5	0.901	0.947
6	0.904	0.949
7	0.900	0.947
8	0.916	0.956
10	0.843	0.912
11	0.855	0.921
12	0.887	0.939
BSP14 average	0.883	0.937
13	0.564	0.708
14	0.688	0.797
15	0.629	0.768
16	0.624	0.765
17	0.757	0.858
18	0.810	0.895
Sonata average	0.679	0.799

The results show that BSP14 plants on average have an 0.88 accuracy and the Sonata plants have a 0.67 accuracy calculated using the Jaccard score. The missing leaves were eliminated earlier on from the automated leaf centre detection or the vein location Sections of the pipeline. This is a beneficial result because the ultimate aim is to automatically select the best leaves for further analysis in Chapter 6. Bright pixels generally indicate that the leaf is well lit and near the light source which would imply the leaf is near the top of the plant. Therefore if the leaf model does not segment all of the leaves due to poor appearance of the leaf, occlusion, etc, this simply filters the available leaf set appropriately earlier on in the process: We would not want to select these leaves for analysis anyway.

The algorithm is not evaluated for the missed leaves, this is not a problem for this case but if the purpose is to segment all of the leaves then the number of leaves missed can be counted from the ground truth and the final segmented leaves. There is ground truth for all of the leaves on the plants and this could determine on a plant level the accuracy for the amount of segmented plant compared to the

ground truth.

4.3 Conclusion

This Chapter focused on building an automated pipeline for segmenting the leaves in a plant. The centre of the leaves are automatically detected and the veins of the leaves are located. These two steps are used to place and align a shape template in the general area of the leaf, then this is incorporated into the level set to find the final leaf boundary/segmentation. Final boundaries and accuracies are encouraging, and poor visible leaves are rejected at various stages in the pipeline.

This pipeline and evaluation used the drought hyperspectral dataset; BSP14 and Sonata plants. This is using the image information in 2D, while this worked well to segment individual leaves the images are hyperspectral data and the plant is imaged in the pot, imaging the plant instead of flat leaves might have an affect on the data. This will be explored in Chapter 5 along with possible solutions if 3D has an affect from imaging plants in the pots (or natural environment).

Chapter 5

3D plant reconstruction and data mapping

5.1 Introduction

Capturing close up images of the plants allows for higher quality data and more subtle changes to be detected, whether that is in the lab, in glasshouses or the field [94]. For hyperspectral imaging the material being imaged for studies is usually flattened onto a surface so the effects of 3D structure are removed - this is not feasible for imaging at high throughput quantities, also it may damage the plant, and it is unnatural. But hyperspectral imaging of 3D objects brings its own challenges. In recent studies it has been demonstrated that the inclinations of the leaves can affect the spectral signatures, and accounting for it can improve the classification accuracy of detecting diseases [10].

There has been a recent review on imaging plants, and in particular the effects of imaging plants in situ [95]. Imaging the plants in 3D rather than flattening the leaf to a 2D surface leads to leaves presenting at different angles and heights. This can change the reflectance [96]. On a slightly larger scale, vertical canopies were imaged to see if there was a difference in the leaf spectral signatures between the upper, middle and lower canopies. There was found to be a difference in several sections of the spectrum range (400-2200nm) [97]. Another study found a difference when the position of the leaves changed, relative to the lighting and the camera. In this case only specific hyperspectral wavelengths were affected [98].

Also, overlapping leaves would mean some leaves are in a shadow, which can also

affect the spectral signatures [99]. These are challenges that will need to be taken into consideration when imaging whole plants in situ if good quality hyperspectral measures are desired.

During this project the plants are imaged in pots. This then includes the problem of imaging the plant with leaves at angles and different heights from the camera (See Fig 5.1). Objects in a shadow are known to have an altered spectral signature compared to the same object that is not covered [99]. If the leaf is in the shadow of another leaf then this could affect whether the plant is considered healthy or diseased. One solution is to disregard occluded and shaded leaves by using approximate depth information to find the height of the plant and individual leaves. A second solution is to measure the brightness of leaves, and have a threshold value which if the pixels' brightness decreases lower than this value the leaves are removed from the analysis. This would need to be a variable threshold for each hyperspectral data set because the height of the leaves and the white balance calibration could change between the plants. A third option is to create an accurate 3D model of the plants to determine the location and position of leaves, and account for 3D variability when using hyperspectral data for analysis. This is the approach ultimately developed in this Chapter.

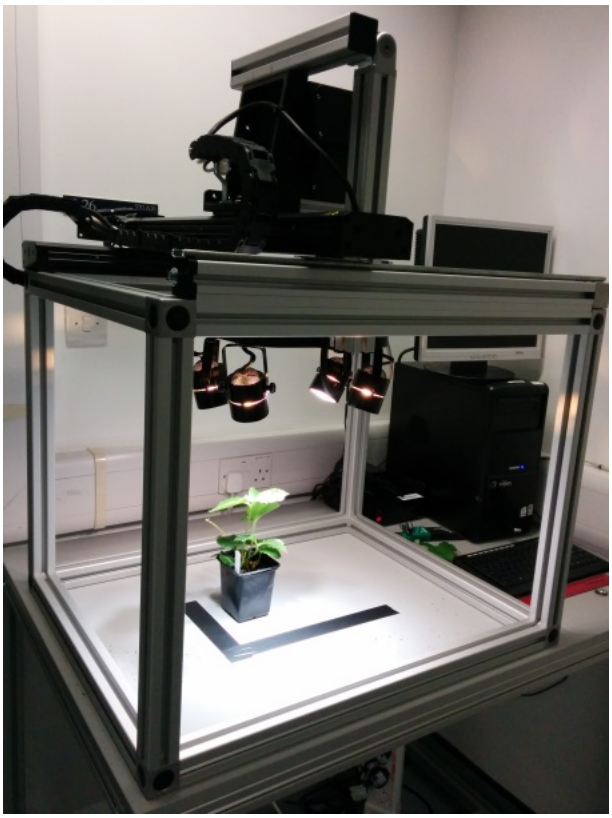

Figure 5.1: A plant in the hyperspectral camera setup. The leaves are at different positions relative to the camera. Flattening leaves is not possible; and this may additionally stress the plant.

The first Section of this Chapter looks at methods that can be used to find 3D information from the hyperspectral data. The hyperspectral data was captured from a top down view of the plant and a pair of hyperspectral images were captured for each plant with a 10cm offset between the two scans. At the same time 3-10 images were captured with a SLR camera from the side of the plant. This data would allow for either stereo pairs or multi-view 3D reconstruction methods to be used, and both are investigated here.

This leads on to the second Section which focuses on mapping the hyperspectral data onto 3D surface models. By the end of the Chapter, an approach will produce 3D hyperspectral plant models which can be used to find leaf location and geometry information, and provide criteria to help select the best leaves for analysis - this is then considered in Chapter 6.

5.2 Imaging plants in situ and the 3D effects on spectral signatures

Imaging plants close up and in situ involves lighting effects, shadows and leaves at different angles. The lights for the hyperspectral imaging setup in Nottingham were attached to a metal frame and pointing towards the center where the plants were located. If the plant was large then the light decreases significantly near the edge of the plant, and if the plant was small the light covers the plant in a more uniform fashion. The best light coverage was in the middle of the image. The upper leaves cast shadows on the lower leaves because the lights were next to the camera and oriented down. Therefore, the top leaves of the plant would be well-lit, but lower leaves might have shadows or be partially obscured from the camera. Finally, as literature shows, the leaf angles can affect the spectral signatures.

There have been studies that have looked at the leaf geometry to see how it can change the spectral signatures of the leaves (described in the introduction). Also an example is now presented to see whether the 3D properties should be further explored in the rest of the thesis.

5.2.1 An example of the leaf orientation influencing the spectral signature

Figure 5.4 displays the spectral signatures of leaves at different angles relative to the camera.

Figure 5.2a is a top down view of the plant; a lot of the leaves appear relatively flat. However they have varying degrees of inclines from the camera plane. Figure 5.2b is a corresponding image with labelled leaves (manually annotated ground truth). Leaves 17,18 and 19 spectral signatures have also been plotted. Leaf 17 has an extreme inclination towards the camera. The reflectance profile for leaf 17 is different to versus other profiles - it is increased in the blue and red wavelengths of the visible light (400-500nm and 570-680nm) while decreasing lower than leaf 19 in the green wavelengths (500-570nm) and both profiles in the near-infrared wavelengths (700-900nm). Therefore, even angle alone can be demonstrated to have an effect on reflectance information.

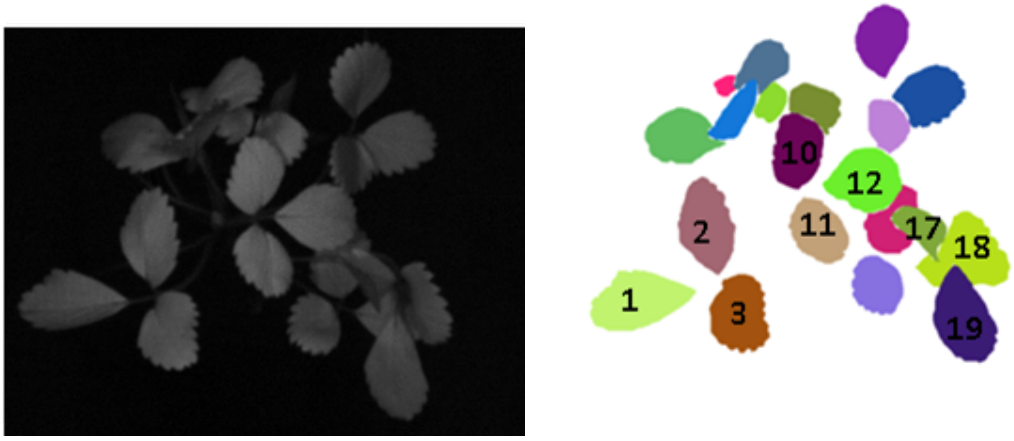

(a) Hyperspectral image - displayed as a grayscale image.

(b) Labelled data - ground truth.

Figure 5.2: Hyperspectral data and labelled leaf numbers.

(a) Side SLR view of the same plant as in Fig. 5.2.

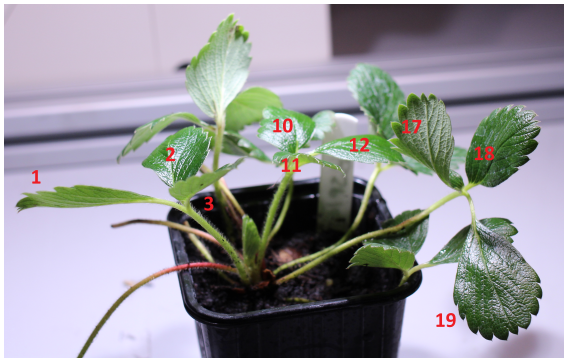

(b) SLR image with leaf numbers that correspond to the leaf numbers in 5.2b.

Figure 5.3: Two images captured from the side of the plant from Figure 5.2.(a) is a side view image of the plant. (b) a similar side view image with numbers that correspond to the leaf ground truth 5.2b. The leaves are at different angles to the camera (top down view) however, the leaves in the hyperspectral image alone (5.2a) appear flat and it is hard to tell the leaf shape from this data alone. Leaves 18 and 19 have a consistent sloping angle however leaf 17 has a much steeper angle to the camera.

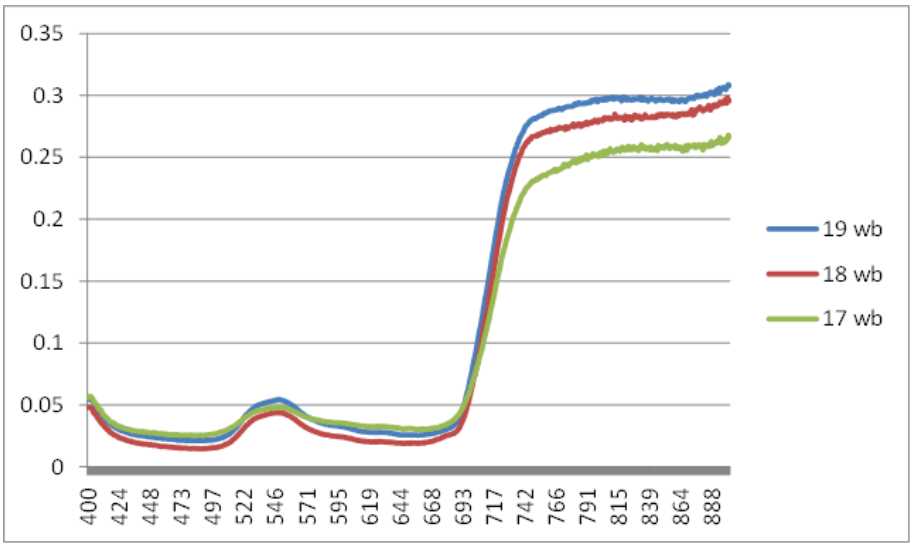

Figure 5.4: Graph of hyperspectral data from three leaves. Leaf 17 is at a steep incline. The profile for leaf 17 is different compared to leaves 18 and 19. The reflectance has decreased in the green and near-infrared ranges and increased in the blue and red wavelengths.

5.3 Overview of 3D reconstruction

There are many existing methods that can produce a 3D model of plants. [100] explores the different broad approaches for 3D reconstruction of plants. There are two main categories: Rule-based methods and image-based methods. The rule-based method incorporates user defined rules to generate example 3D models. These are synthetic plant structures developed via a modelling process, and do not capture the detailed structure of real plants. This method is good to visualise and research how the plants respond to environmental changes. The second category is that of image-based methods which incorporate computer vision techniques and extract the geometric information from the plant data collected to reconstruct accurate 3D models.

Image-based methods can themselves be split into two groups: active and passive. Both methods require hardware. Active approaches use controlled light. Passive approaches only require a camera that captures overlapping images of the plant. The passive image based method is more accessible and the option that is going to be discussed in more detail here.

There are several approaches that use different properties of the images; Shape from silhouette [101] views the silhouettes from the camera's viewpoints to build

the 3D model. Space carving [102] involves a bounding box of the object which is separated into groups of voxels, and if the voxels are consistent between the images then the voxels are included in the final model.

Binocular stereo vision captures two images on the same plane focused on the object with an offset between the left and right image. Then the depth map can be created by matching the features (edges, lines, pixel intensities) in the two images, measure the distance and the distance values can create a disparity map. This is based on human vision, where the eyes view two different viewpoints with a small offset. A close by object has a bigger difference between the left and right eye compared to a far away object where the difference can be hard to notice because the difference is small.

First the cameras must be calibrated, this is finding the distance between the two cameras and the rotation angle between the cameras pointing at the object. The rotation matrix and translation vector need to be found. Then the images need to be aligned to account for the different viewpoints and distance from the object, this aligns the two images horizontally. The final step is stereo matching, this can use local or global methods. Local methods use block matching, feature matching, etc whereas global methods use intrinsic curves, graph cuts, etc [103]. Local methods can mismatch if the images are noisy/low resolution, while global methods have a longer computation time.

Multi view stereo extends to multiple images, captured at different view points of the object to build a precise 3D reconstruction model. Structure from motion (SFM) estimates the images location and the camera parameters by matching the features across images and locate where the cameras are in relation to the matched features [104]; the overview is: Detect 2D features in every image; these could be edges, points, lines, regions and they could be detected using Surf [105] or Scale invarient feature transform [106] as two examples. SIFT extracts features and identifies them with a unique feature description so they can be found regardless of scale or illumination, then the Euclidean distance can be used to find the similarity between the features [107]. The matched features, correspondence points, can be used to calculate where the cameras are by using the features to project tracks to the same point in space (3D). The more features that match and the more images that are included adjusts the cameras to have a more accurate position for each camera and builds up a point cloud in 3D.

Multiview uses the camera parameters and adjusted image locations to project matched features (correspondance points mentioned previously) and where the

two projections interesct this should be the 3D point [108].

5.4 Stereo pair hyperspectral data

Stereo pair hyperspectral data was collected for the drought experiment. After the first hyperspectral image was captured the camera was moved 10cm to the right and imaged again (the camera specifications are listed in Chapter 3).

A pair of images were then captured to create a depth map. All of the plants for drought were imaged twice to capture the stereo pair data. Here Sonata and BSP14 are again considered; this includes 18 plants.

The hyperspectral pairs were captured at the same x start and end positions, with the y start position moved by 10cm for the second image; therefore the corresponding pixels between the images should be on the same row, within a tolerance. The hyperspectral image pair are reduced to a grayscale image pair at 750nm to make it easier to match small pixel blocks between the images.

For each block (9x9pixels) in the left image the sum of absolute differences (SAD) is calculated between the block and pixel blocks along the same row of the right image. The block with the smallest SAD is selected as the the correct matching blocks and the distance between the two blocks is the disparity value. Repeating this process for all the blocks in the left image results in a disparity matrix (disparity map) where the lowest value is the closest point to the camera and the highest value is the furthest point away from the camera [109].

Figure 5.5 has an RGB image of the plant to visualise the plant and the disparity map where the light grey pixels are higher on the plant (closer to the camera) than the dark grey pixels.

(a) Colour image of BSP14 - control plant 12 to visualise the plant

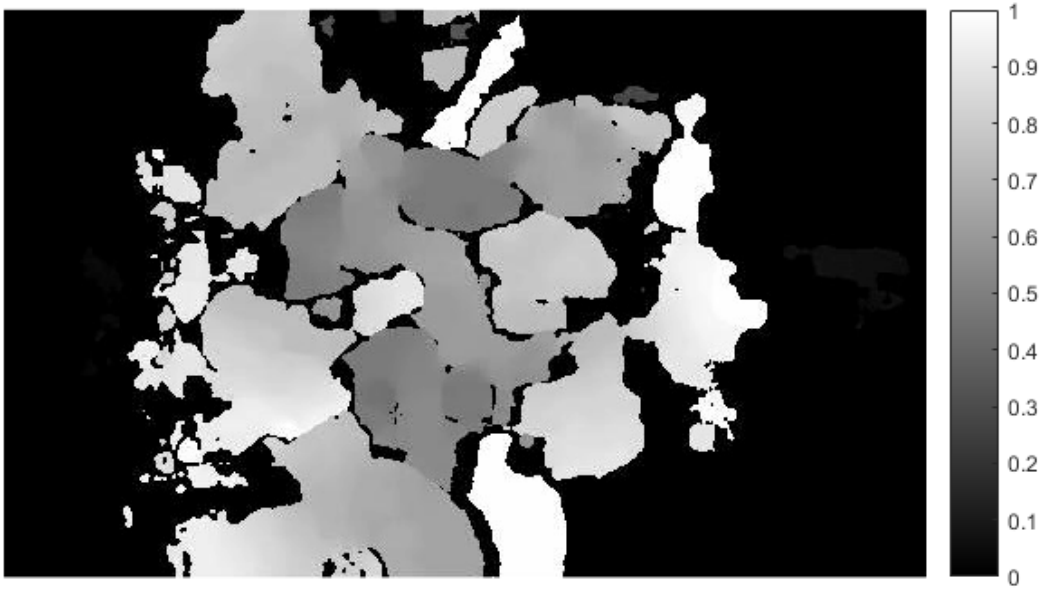

(b) Disparity map of the whole image.

Figure 5.5: Comparison of the colour image and the disparity (depth) map. (a) is the RGB images created from the hyperspectral data and (b) is the disparity map (displayed as a colourmap for clarity). The disparity map is quite noisy however there is a slight difference in areas of the plant image which in general is indicating the distance of the leaves from the sensor.

hyperspectral data is too low resolution to enable a good correspondence to be calculated.

5.5 3D reconstruction and mapping hyperspectral data to the 3D model

This Chapter so far has looked at the effects of imaging plants in situ and how the leaf angles can change the hyperspectral reflectance profiles. A previous study showed that shadows can also influence the profiles [99]. With the leaf angles, shadows and lighting all influencing the hyperspectral data this needs to be taken into consideration. Therefore, at least the leaf angles and height need to be calculated in order to discount leaves that may lead to poor measurements. Stereo methods solely using the hyperspectral camera could not determine the leaf angle or the height. Therefore full 3D plant models will be needed and the hyperspectral data will need to be mapped to the 3D model.

The plan is to have hyperspectral data registered onto a 3D model of the plant in order to obtain the orientation, angle and location of the leaves. Multi-view images were captured at the same time as the hyperspectral data, this means a model can be built using a multi-view stereo reconstruction method and the hyperspectral data later mapped on. There are a few methods currently available for 3D reconstruction that use multiple views discussed at the beginning of the Chapter. Patch-based multi view stereo approaches are an option that are used to produce a point cloud in 3D space.

The pipeline is outlined in Figure 5.6. First the plants are reconstructed using the 3-10 SLR side view colour images captured at the same time as the hyperspectral data. The method used is PMVS which is a multi-view stereo method and this creates a dense point cloud of the plant. Then the surface reconstruction software is used to create a 3D model for the plant. This is a colourless surface model with vertices and faces. The third step involves projecting the 3D model to 2D and finally the fourth step is to register the hyperspectral image to the top down view of the 3D model followed by mapping the hyperspectral data onto the 3D model. These patches with mapped hyperspectral data are referred to as hyperpatches. A flow-chart detailing their creation is shown Figure 5.6.

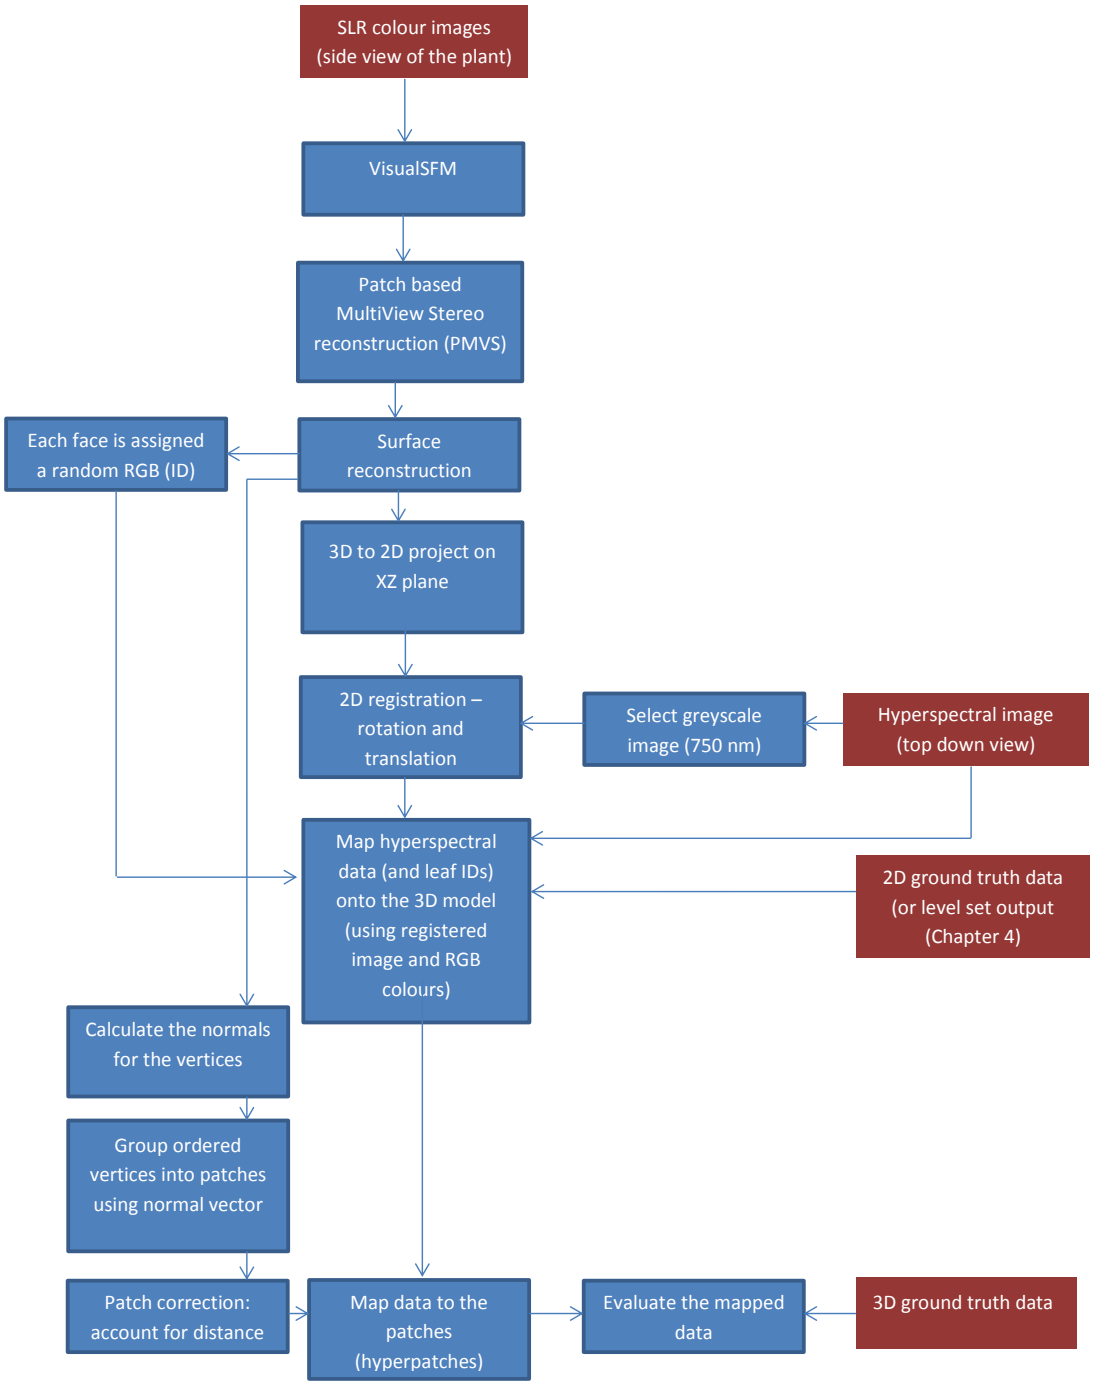

Figure 5.6: Pipeline for the 3D reconstruction, hyperspectral data mapping and hyperpatches.

5.5.1 Patch based multi view stereo software

First of all, to use Patch based multi view stereo software (PMVS) the camera parameters need to be known [110]. To do this VisualSFM can be used to find the camera parameters using structure from motion (SFM) [111] [112]. The same basic principles are used as for stereo vision. This process matches up corresponding features in the images to locate the camera positions. The features used are Harris and Difference-of-Gaussian operators which detect corners and blobs. Once the features are matched the patches are tested for photo consistency using normalised cross correlation - checking the consistency of the colour in the voxels/patches.

PMVS takes the RGB SLR images (Fig. 5.7) and camera parameters as inputs to build a 3D point cloud where the colour information can be included on the points. The output is a noisy point cloud which contains patches of the 3D model of the plant. However at this point the model has large sections missing, especially from a view point not included in the original images (See Figure (5.8).

Figure 5.7: Original images used for PMVS and surface reconstruction. These are SLR colour images captured from the side of the plant. 3-10 images were captured for each plant and at each imaging session for the drought dataset captured in Nottingham.

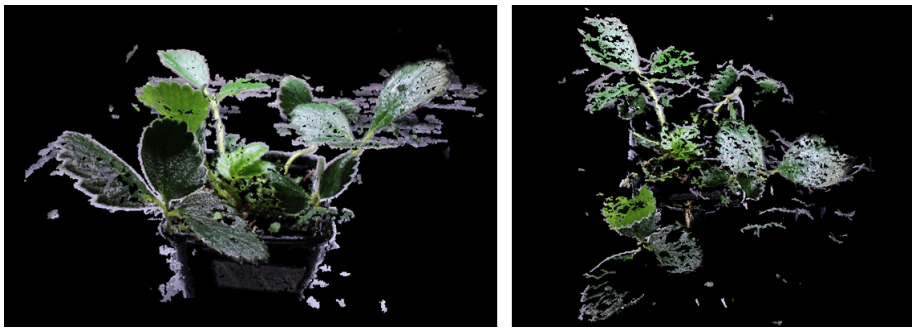

Figure 5.8: Patch based 3D reconstruction including the colour texture. Left: Side view of the model where most of the images were captured from. Right: Top down view displaying gaps where no features were detected due to occlusion and view restriction.

5.5.2 Surface reconstruction

Next surface reconstruction software is used. This software was created for the reconstruction of plant shoots [113]. The surface reconstruction software takes the images, the dense point cloud from PMVS and camera parameters as inputs to reconstruct a 3D surface model.

The software uses the output from PMVS, and the dense point cloud, and represents the surfaces of the plant as small patches. The patches are placed using the PMVS output then they are adjusted using level sets to connect and align neighbouring surfaces. The patches themselves consist of a group of triangles (faces) and each triangle (face) contains three vertices *(x1,y1,z1), (x2,y2,z2)* and *(x3,y3,z3)*. There can still be spaces in between patches due to the parameter settings of the software if not enough images were captured; this is not a problem for this application.

Figure 5.9 shows the result of the surface reconstruction software on one plant. The surfaces resemble the plant; however there is no longer colour information on the model at this stage.

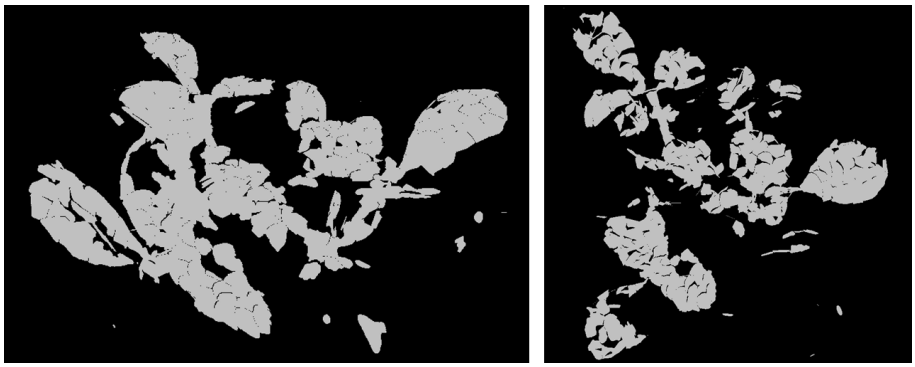

Figure 5.9: Surface reconstruction has created patches that has improved the leaves where information was missing. It also smooths the patches, however the colour and texture is missing from the original images.

5.5.3 Registering the hyperspectral data

The hyperspectral data is captured via one top down view of the plant, and from the 3D reconstruction there is now a corresponding 3D surface model of the same plant as is the hyperspectral image. The hyperspectral data and the 3D model were captured and created using different sensors and at different heights and angles. Therefore they will need to be aligned as part of a registration process before mapping of hyperspectral data onto the 3D model can occur.

Registration in 2D

The registration process here is in two parts; the gross orientation of the 3D model (to find the top down view that the hyperspectral images were captured at), and the fine tuning registration.

Gross orientation - global registration

The hyperspectral data is a top down view of the plant in 2D. The first step of the registration is to find the view that best matches the 2D hyperspectral data by rotating the 3D model and projecting the view to a 2D plane. The initial top down view (orientation) has been manually selected because the camera setup is known therefore the known geometry can give an approximate initial alignment. Once the orientation was found for one plant it worked for the whole dataset. Once the general orientation was selected it could be used for the whole dataset therefore the manual selection has been applied using the known geometry for this part.

Fine tuning - local registration

The 2D view of the 3D model that best matches the hyperspectral image is selected and projected, in the case of this dataset the camera is on the x,z axis with the camera facing along the y axis. The hyperspectral image is scaled to be a similar size as the projected view by locating the furthest x and y pixels on the plant and the projected model (x and z) and then scale up the image by the distance between the points. An automated registration process then uses a grayscale 2D projection of the model and a grayscale image of the hyperspectral data.

The projected model and greyscale (hyperspectral) image are different, the model has no texture and the general shape of the leaf edges will appear but there might be gaps where the patches join. The hyperspectral image is low resolution and noisy. There are images that are similar but have been captured using different imaging systems or from different time points, images can change over a time series or between the imaging devices however the images need to be compared and one popular metthod used for medical image registration is mutual information. This is when features need to be matched between the images but the features have changed slightly, therefore the statistical information is compared to find the closest match [114]. The greyscale hyperspectral image is translated and rotated until the mutual information is maximised between the projected 2D model and the hyperspectral image.

The 2D registration results are displayed in Figures 5.10 and 5.11. These figures show the two error modes of registration.

Figure 5.10 shows the missing 3D surface data. The images used to create the 3D model are at a restricted side view angle and with a limited number of images. At the time of collecting the images it was not known that the images would be used for 3D reconstruction. They were collected knowing 3D information was needed and multi-view reconstruction is possible with a limited number of images. It was also hard to capture images directly over head when the plants were in the camera setup and it was best to maintain the plants position for the hyperspectral images and the SLR images. This is why some of the 3D model is missing when the plant model is viewed from the point of view of the hyperspectral camera (directly above), and could be improved by collecting more images next time.

Figure (5.11) displays the hyperspectral data in grey pixels on the 2D version of the plant model. The white sections are missing hyperspectral data. There are a few possible reasons why the hyperspectral data is missing. 1) misregistration: a) the 2D projection of the 3D model is not exactly aligned with the y axis. b) the translation is slightly wrong. 2)extra triangles/patches reconstructed on the 3D

model or noise.

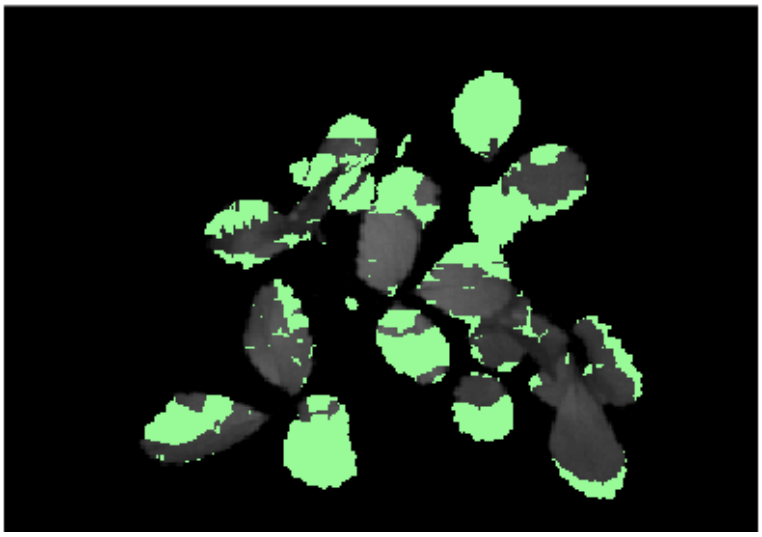

Figure 5.10: 2D registration - The grey pixels are the 2D hyperspectral image (displayed at 600nm) and the green pixels represent missing 3D model for the top down orientation. 2D hyperspectral data exists but no 3D surface model exists. The 3D sections (on the 2D projection) are missing because the 3D reconstruction only had images from the side of the plant and this is the top down view of the plant that the hyperspectral images were captured at.

Figure 5.11: The grey pixels are hyperspectral data and the pink pixels represent missing hyperspectral data there is a 3D surface model (on 2D projection) but no hyperspectral data (that could be mapped onto the 3D surface model).

5.5.4 Mapping hyperspectral data onto the 3D surface model

The next step is to map the hyperspectral data onto the 3D model. The model is at this point a collection of patches representing leaves and within the patches are a collection of triangles. The triangles in the 3D model are randomly allocated RGB values as ID numbers which are unique (Figure 5.12); the RGB combinations are compared and any repeating combinations are replaced until the list is unique. The 2D top-down projection of the 3D model has been registered with the hyperspectral data. Therefore, the hyperspectral image (x,y) is aligned with the 3D model (x,z). Then for each triangle of the 3D model from the top down view the registered hyperspectral data is selected. One triangle does not directly equal one pixel from the hyperspectral image. Sometimes for one triangle there will be several 2D hyperspectral pixels, and the average pixel value for the triangle is calculated and mapped onto the model. If the triangles are smaller than the size of the hyperspectral pixel then the pixel is mapped onto the set of triangles that are within the pixel area. Finally, there will be random triangles that do not have hyperspectral data for a few possible reasons: the top down view obscures the leaf underneath, misregistration, missing hyperspectral data due to the leaf being cut off in the image (this is displayed in pink in the following Figures).

The hyperspectral data mapped to the corresponding location on the 3D model can be visualised (Figure 5.13). There are a few problems with the current process. If the registration is slightly misaligned then the 2D leaves will be mapped onto the wrong 3D leaf. Figure 5.16 displays the individual leaves mapped onto the model; the bright green is partially projected onto the leaf below, which should be labelled as missing data. Also, the red sections show patches that due to the angle of the leaves mean they are not presented in the top down view.

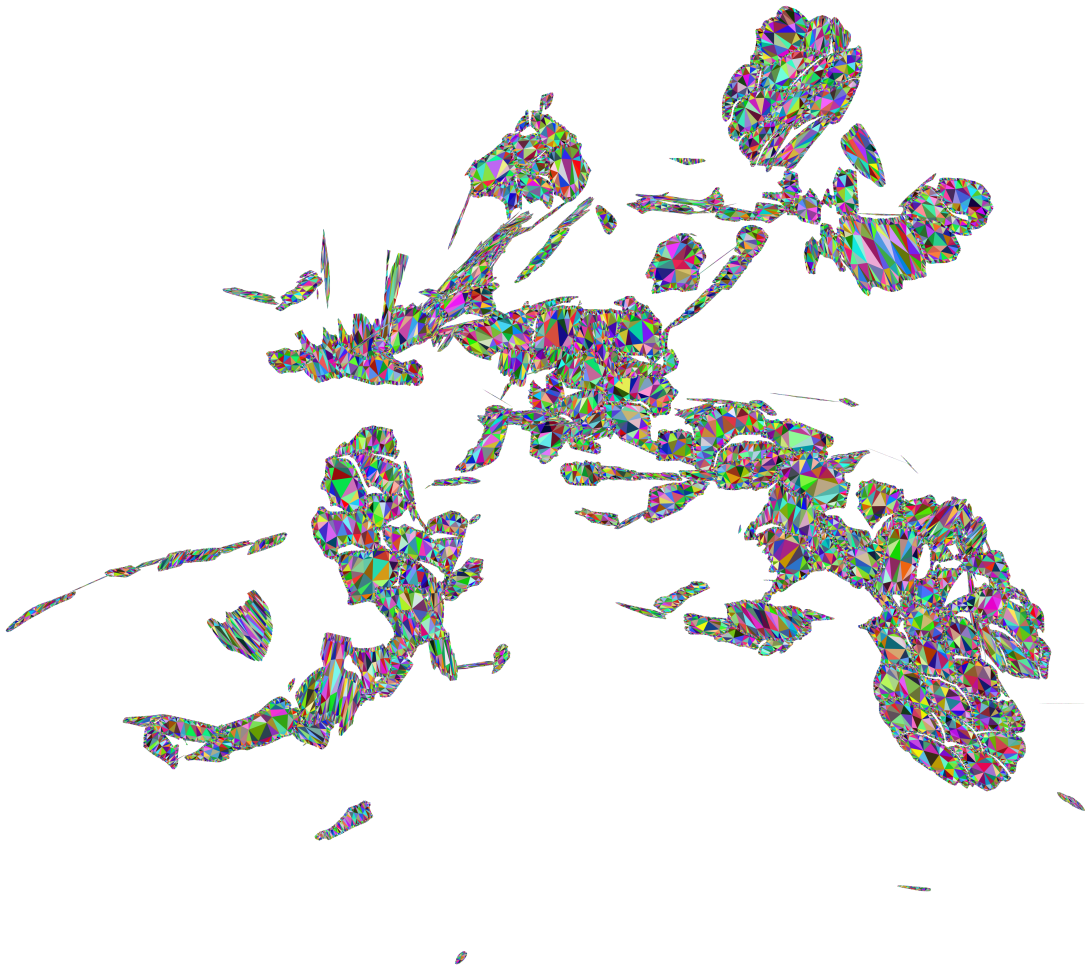

Figure 5.12: The 2D top down view of the 3D model, showing unique ID values for triangles respresented at RGB colours.

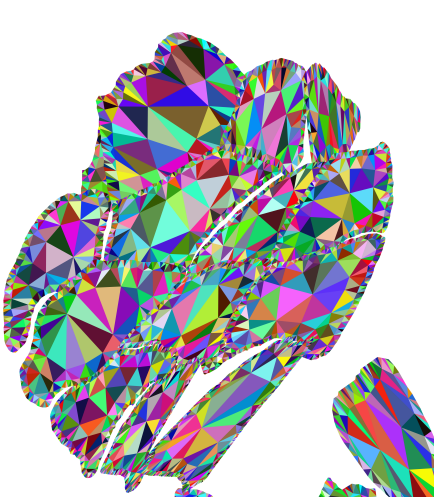

(a) Top leaf of BSP14 plant 11 from the 3D model with the random colour triangles.

(b) A leaf from BSP14 plant 11. The hyperspectral data is mapped onto the 3D model (dark grey) and the pink is the missed hyperspectral data where a model exists but the data does not.

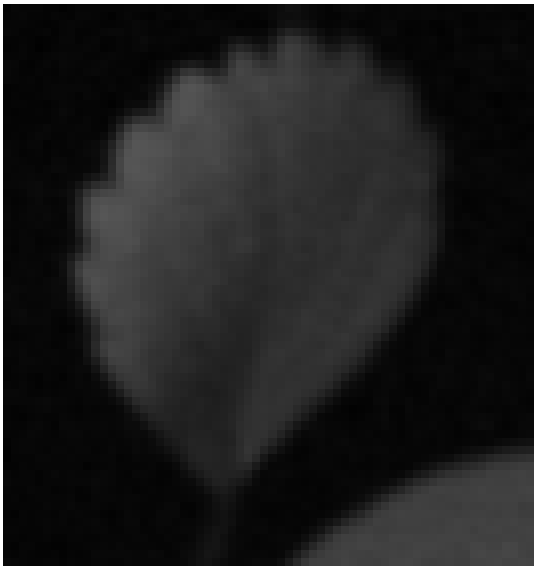

(c) A grayscale image of the same leaf from the hyperspectral image.

(d) The same leaf from the slr colour image.

Figure 5.13: A leaf from BSP14 plant 11 to display what the leaf looks like in different modalities. (a) 3D model of the leaf with the random triangles. (b) the same model leaf with mapped hyperspectral data (displayed using one wavelength) (c) The corresponding leaf taken from the hyperspectral image and displayed as a grayscale image. (c) the same leaf again taken from a colour image.

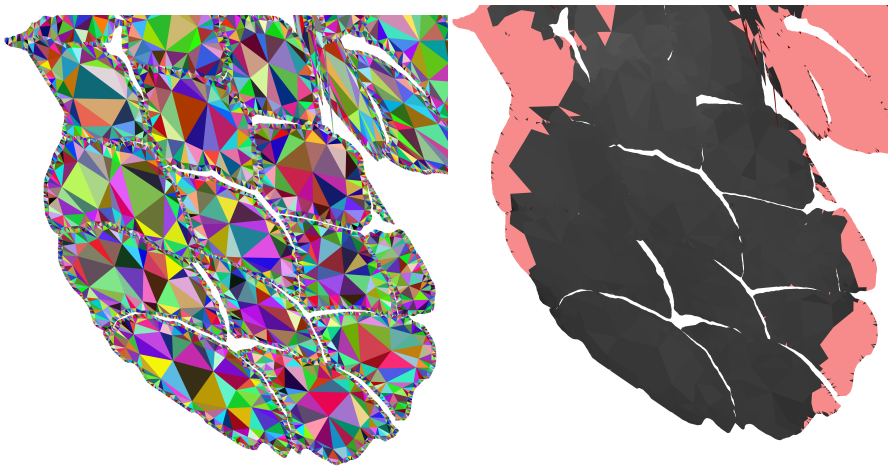

(a) A second leaf on BSP14 plant 11. This is a leaf from the 3D model.

(b) A second leaf from BSP14 plant 11. This is the hyperspectral data mapped onto the leaf (dark grey) and the pink is where the hyperspectral data could not be mapped (the missed data from the 2D registration).

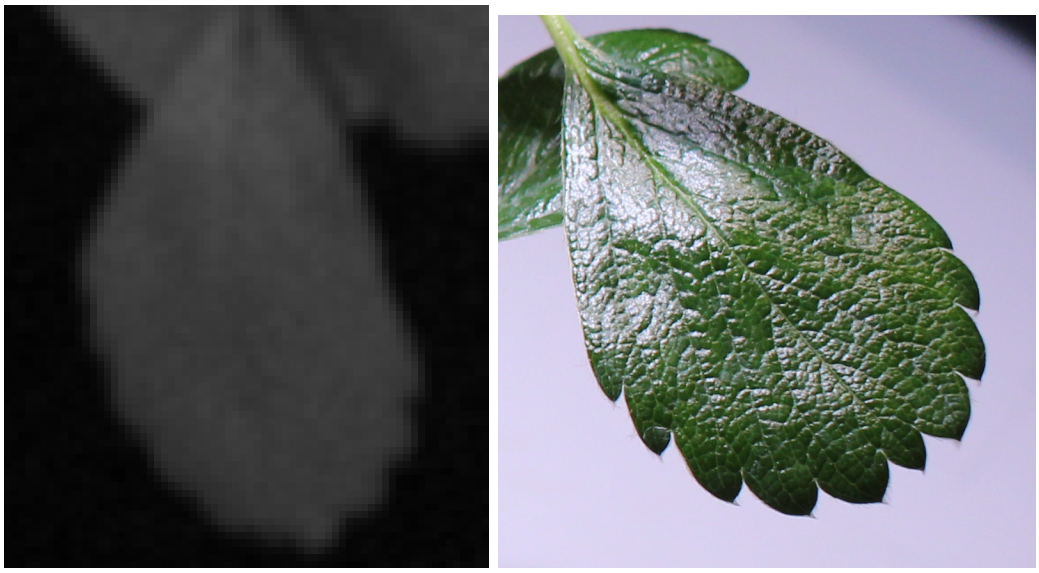

(c) The same leaf from the hyperspectral image.

(d) Corresponding leaf taken from a RGB SLR image.

Figure 5.14: A second leaf from BSP14 plant 11. The 3D model leaf, the mapped hyperspectral data, the grayscale hyperspectral leaf and the colour leaf images are displayed together.

The hyperspectral data (Fig. 5.15b) and the 2D ground truth, also called leaf IDs (Fig. 5.15a), are mapped on to the 3D model. The hyperspectral data is mapped to obtain 3D hyperspectral data and the leaf IDs are mapped to know where the hyperspectral data is mapped onto the 3D model. The leaf IDs are also easier to

visualise, and check if the mapping worked. This is used to evaluate how much of each hyperspectral leaf data is mapped onto the correct leaves in the 3D model.

The leaf IDs and the hyperspectral data have been mapped onto the 3D model and can be visualised in Figures 5.16 and 5.17. The leaf ID colours from the ground truth Figure (5.15a) correspond to the leaf colours in the 3D model (Figs. 5.16 , 5.17). The hyperspectral data is harder to interpret but the hyperspectral leaf data and the leaf IDs originate from the same 2D image, therefore the leaf ID model is a good visual indicator to investigate what is happening. The pink sections are also mapped from the registration process. These sections are where the hyperspectral data was missing from the top down view of the 3D model. The red sections are where the hyperspectral data was missing because the hyperspectral images were captured from above the plant and the leaves below would be occluded or leaves at a steep angle would miss some of the hyperspectral data being mapped onto the model. The red sections are more prominent in the side view of the 3D model (Fig. 5.16) than the top down view (Fig. 5.17) where the red sections should be minimal.

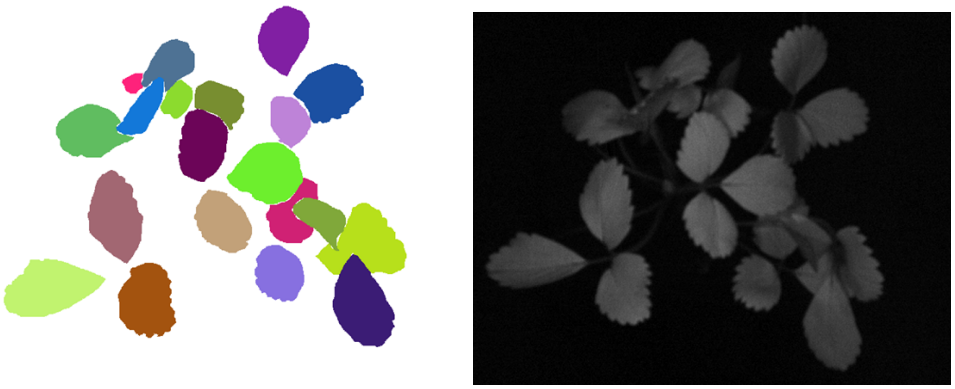

(a) **(b)**

Figure 5.15: (a) The labelled ground truth image displayed as a colour map to visualise the different leaves. (b)The 2D top down view of the hyperspectral data displayed as a grayscale image

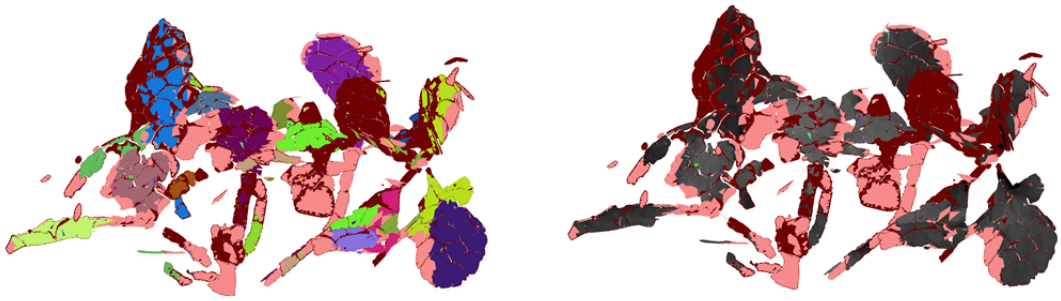

Figure 5.16: Side view of the 3D model. Left: individual leaf IDs mapped onto the 3D model. Right: hyperspectral data (600nm) mapped to the model. The pink highlights where the hyperspectral data is missed from the model. The red is missed due to the angle of the leaves causing problems for the top down view

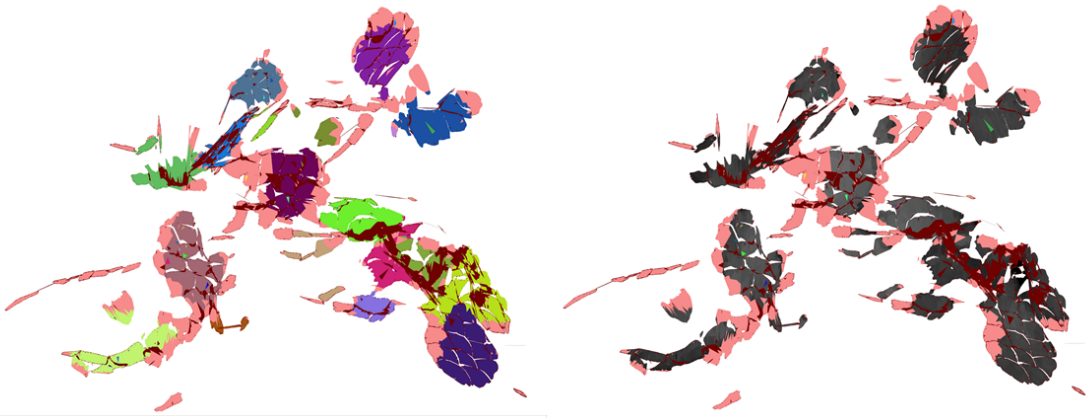

Figure 5.17: Top down view of the 3D model. Left: individual leaf colours mapped onto the model (same colours as Fig. 4). Right: the mapped hyperspectral data (600nm). As before the pink highlights where the hyperspectral data is missed from the model. The red is missed due to the angle of the leaves causing problems for the top down view.

At this stage, then, there exists a pipeline for registering hyperspectral data onto a 3D model. The 3D model has been constructed from multiple side view images. However, the hyperspectral image is a single, top down view. This means there are inevitably small sections of the model missing, and conversely hyperspectral data is missing from lower leaves and leaves that were at an acute angle to the camera.

This process only uses 1 hyperspectral top down view of the plant and 3-10 digital SLR images of the plant from different angles. The processing time per plant is

35 minutes: 5 minutes for PMVS, 25 minutes for surface reconstruction and 5 minutes for the registration.

At this point the hyperspectral data is mapped onto the 3D model per triangle but the patch information that was created in the surface reconstruction software has been lost. The only information retained for the 3D models is the list of triangle vertices and faces.

5.6 Manual creation of 3D ground truth

To evaluate the registration of the hyperspectral data on the 3D model there needs to be ground truth data for the 3D model. To create this, the leaves are manually labelled using different colours in Meshlab [115], then the faces, vertices and colour information are saved in the 3D model. The ground truth data was created for the Sonata and BSP14 plants for day 3 and 4. BSP14 plant 3 and 9 could not be reconstructed and therefore there is no ground truth or analysis for these plants.

An example of the manual selection process is shown in Figure 5.20. Where Figure 5.18 displays two different views of the plant, then Figure 5.19 is the point cloud produced. However Figure 5.20 is the surface reconstruction where all colour is removed and the only information left is the patches in 3D. A few leaves have been coloured however in the lower right section the blue and purple leaves overlap and the patches do not exactly follow the leaf edges.

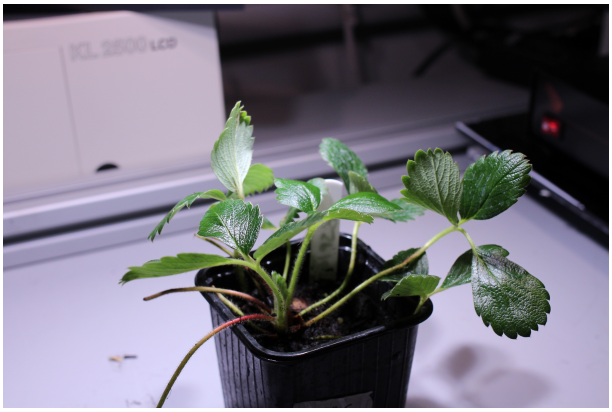

(a) Point of view of BSP14 plant 11.

(b) Second point of view of BSP14 plant 11

Figure 5.18: Two views of BSP14 plant 11. The images were captured using a D-SLR camera just before the hyperspectral imaging occured.

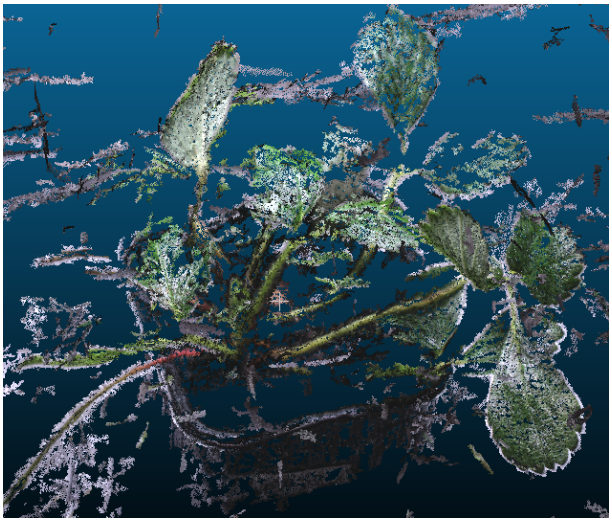

Figure 5.19: Point cloud of BSP14 plant 11 to display the colour and details of the plant for a comparison when labelling the 3D model.

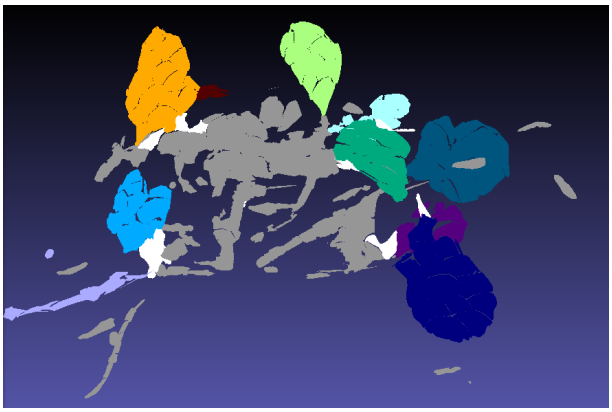

Figure 5.20: BSP14 Plant 11. A few leaves have been manually selected and display the ground truth for the leaf using colour.

Just for display here two plants have been manually selected for the ground truth (BSP14 Plants 11 and 12). The three Figures (5.21 - 5.23) display different view points of the 3D model for plant 12. White is the colour when the patches have not been assigned to a leaf, this will contain the stem, stolon, noise and any leaves or patches that were hard to determine if they belonged to a leaf.

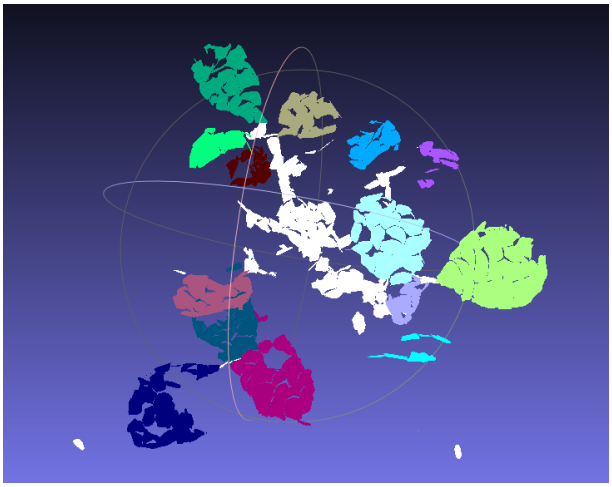

Figure 5.21: BSP14 plant 12 view 1. This is a top down view of the plant, the same view point that the hyperspectral images were captured.

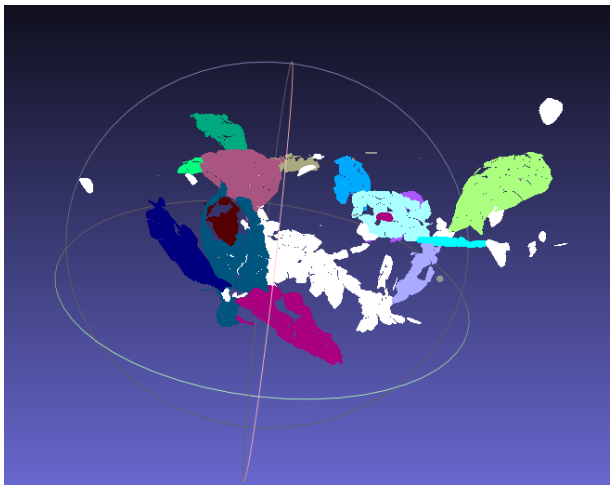

Figure 5.22: BSP14 plant 12 view 2. This is a side view of the plant, the view point where some SLR images were captured from that were used to create the 3D model.

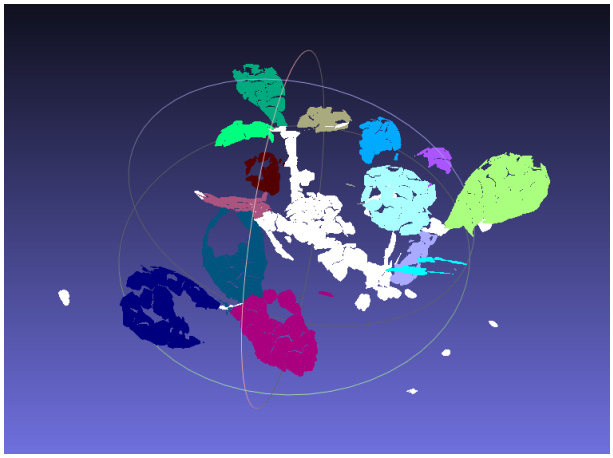

Figure 5.23: BSP14 plant 12 view 3. This is view in between the top down and the side views.

5.7 Results and comparison of the ground truth and mapped data

The mapped hyperspectral data (equivalently, the leaf IDs) are compared with the 3D manually labelled ground truth data to look at the quality of the mapping from the 2D source image to the 3D model. Figure 5.24 displays the ground truth and mapped leaf IDs for BSP14 plant 12. The colours from the ground truth and mapped leaf IDs are not the same. The data is being mapped to the triangles, not the patches, in this Section.

(a) Annotated 3D model of Plant 12 (Ground truth)

(b) Mapped model of Plant 12, pink is the sections where there is no hyperspectral data mapped onto the 3D model.

Figure 5.24: The ground truth model and mapped model in 3D, from a top down view of the plant.

Due to the process of mapping the data to the 3D model the upper most visible leaves are mapped, and any hidden lower leaves will of course be missed. The hyperspectral data is captured from a top down view of the plant which means the leaf data underneath the top leaves is missed. However because there is likely high quality hyperspectral data for the top layer leaves it doesn't matter if lower levels are missed in the model.

Figure 5.25 displays the effect of two overlapping leaves: (a, c) are the ground truth leaves and (b, d) are the mapped leaf IDs. This displays the top down view of the 3D model and leaves 9 and 15 appear fine (b, d) compared to the side view.

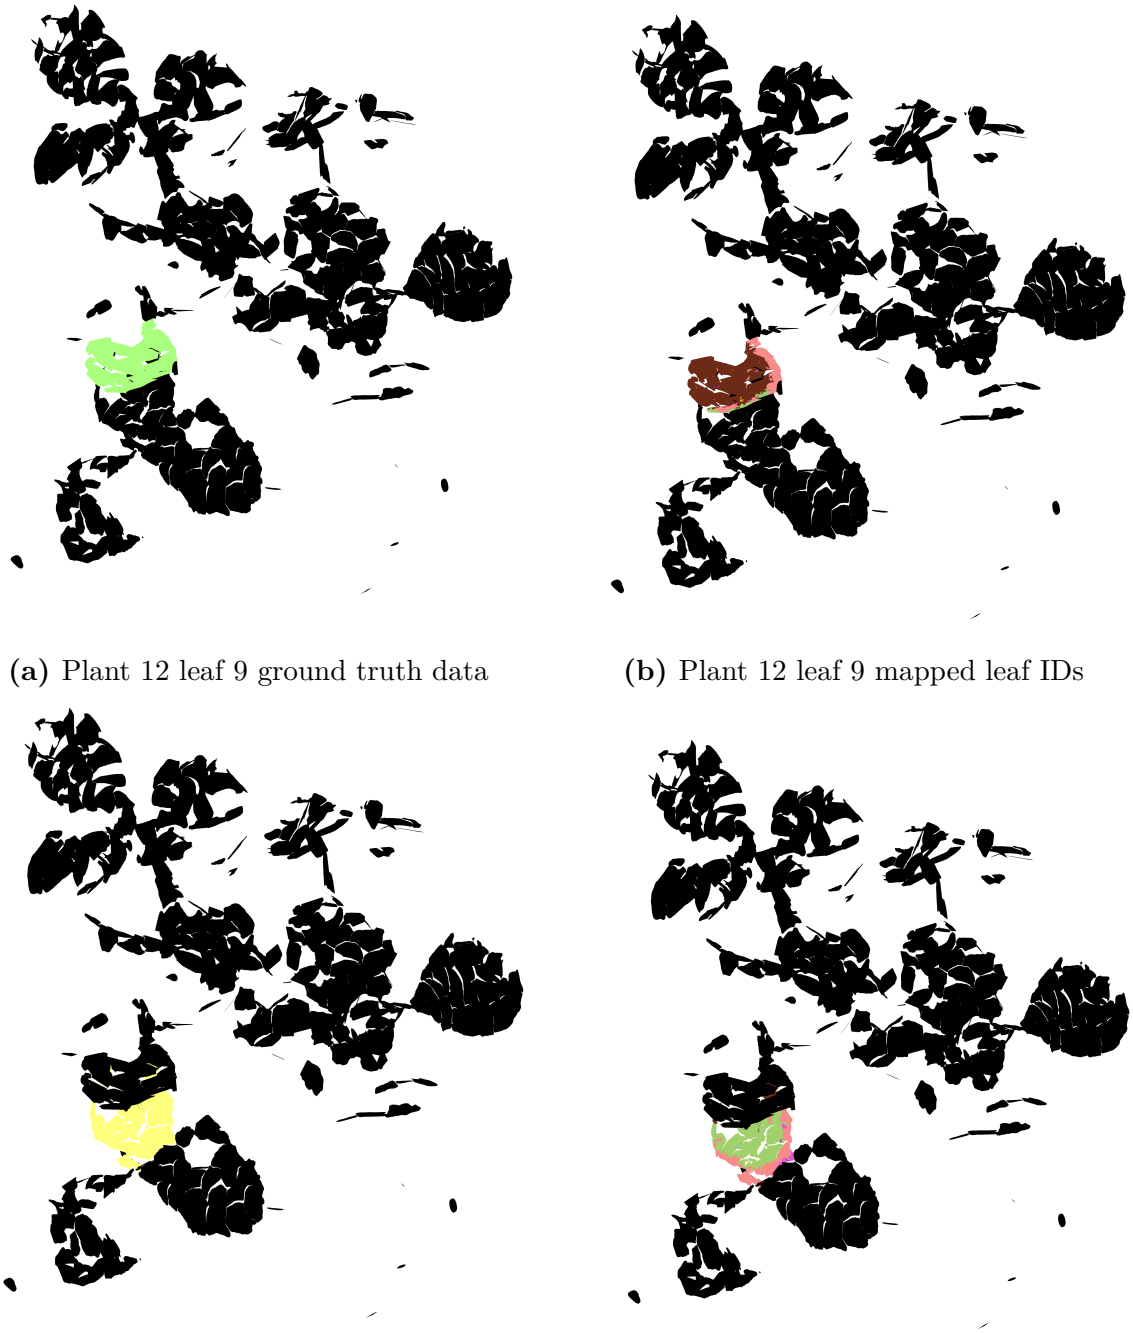

(a) Plant 12 leaf 9 ground truth data

(b) Plant 12 leaf 9 mapped leaf IDs

(c) Plant 12 leaf 15 ground truth data

(d) Plant 12 leaf 15 mapped leaf IDs

Figure 5.25: BSP14 Plant 12 leaves 9 and 15 comparison. Leaf 9 is overlapping leaf 15, and from the top down view this looks as though it is not a problem however from a side view the leaf ID is not mapped that well due to occlusion. This can be viewed more clearly in the side view in Figure 5.26.

A side view reveals both missed data and misregistration (see Figure 5.26b). Leaf 15 is shown in green and pink is the known missed hyperspectral data. However there is a section of brown from leaf 9 visible on leaf 15, this would count as

mislabelled.

(a) Plant 12 leaf 15 gt side view.

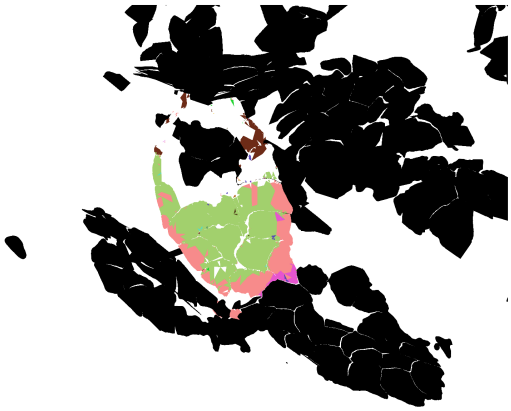

(b) Plant 12 leaf 15 map side view.

Figure 5.26: BSP14 Plant 12 leaf 15 side view comparison. (a) is the ground truth labelled data and (b) is the mapped leaf ID data. The mapped data (green) has a few mislabelled faces. This is a restricted viewpoint for the hyperspectral data and some of the leaf was not visible from the top down view. Therefore only approx half of the leaf is correctly labelled. Pink is the missed hyperspectral data and brown is the ID for leaf 9 which is mislabelled.

5.8 Hyperpatches

The Figures above have had data mapped on a per-triangle rather than per-patch basis. For plant 12, as an example, there are 71k vertices and 90k triangle faces. Also contained in the data there is patch information describing how groups of triangles are collected together for each small patch surface. The surface reconstruction software [113] originally collected vertices into patches however only the vertices and faces data structure was passed through to the 3D model, the patch information has been lost. Developed here is a method to find the patches again by calculating the normals for the triangles. The vertices are in order for the patch information from the surface reconstruction software. This is helpful because it means the patches can be easily selected from the vertices by looking for the change between normals.

The process for reforming patches is as follows:

1. Calculate the normals for the vertices
2. Iterate through the vertices in order (extracted from the surface reconstruction software)
3. When the normals change between two vertices save that group of vertices as one patch
4. Repeat the process until all the vertices have been grouped together into patches.

An example of the patches is displayed in Figure 5.27b.

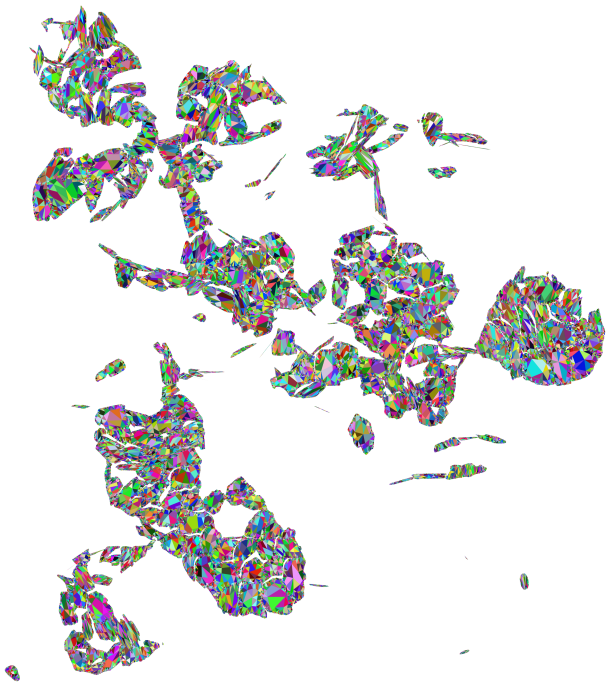

(a) The triangles are assigned a unique random colour.

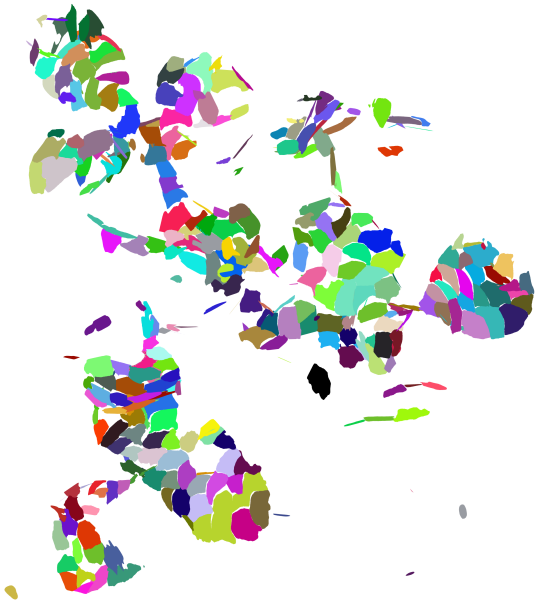

(b) Patches created by combining triangles with the same vertices normals. Random individual colours display the patches. There are small patches combined together where the patches have shared the same normals therefore they were grouped as one patch. The original surface reconstruction software did not group these together.

Figure 5.27: The triangles and patches displayed together to see the difference in detail and complexity. There are 71k triangles and 312 patches.

A problem with the patches method for grouping the triangles is that if the triangles have matching normals and as long as they follow in the vertices list then they will be combined into one patch. This is a clustering type approach and using the above method means 2-3 patches might be combined. As long as the patches are a part of the same leaf then it is not a problem. Figure 5.27b on a lower leaf displays a long patch of light green colour; this is an example of multiple small patches being combined by using this process.

For many faces in practice the algorithm to group the vertices together using the normals will work. However when patch (i) and patch (i+1) have the same normals but they are on seperate leaves they will be mis-classified as the same patch. This results in certain patches being misclassified, displayed in Figure 5.28 (black circles), and a specific example of the problem is shown in Figure 5.29.

Figure 5.28: Initial face to patch selection. There are patches with the wrong leaf ID, a few of the patches have a black circle to highlight that the patches are wrong.

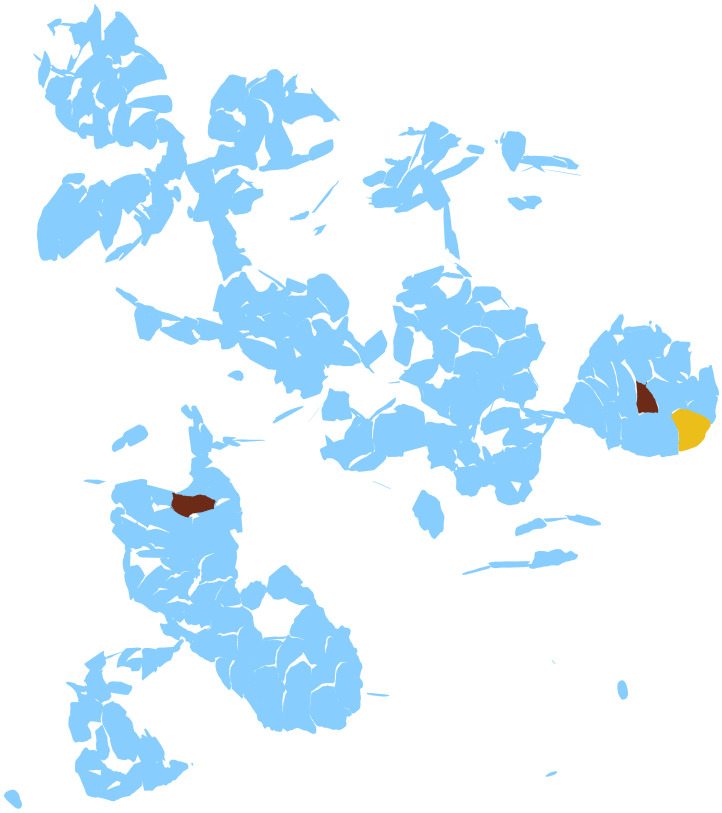

Figure 5.29: Two patches selected (brown and yellow) however the brown patch is over two leaves. Displayed with a light blue background to clearly see the two patches.

One solution to this erroneous grouping across leaves is to take into account extra information known about the model. In this case a distance criteria has been applied. For each patch, the distance of the vertices are measured and if the distance in any direction exceeds a limit then the patch is subdivided. This works for this data, but is sensitive to the distance limit which would most likely need to be amended for significantly different plant datasets. Figure 5.30 is the resulting model after the patch corrections were applied.

Figure 5.30: The patches after accounting for the distance as well as the normals. The leaf IDs are mapped to the patches.

5.8.1 Evaluation of mapping accuracy of the 3D model

For each plant that has mapped hyperspectral data to the patches there is now a manually created 3D ground truth plant (see Section 5.5) where the leaves are labelled from 1 to N (number of leaves). The mapped leaf (leaf ID) is here compared to the corresponding 3D ground truth leaf ID and the number of faces that are correctly mapped onto the leaves are counted.

The manual ground truth accounts for the majority of the leaves on the plant. For BSP14, a few leaves in the middle of the plant were hard to visually determine, therefore they were discounted, and for Sonata two plants were very hard to create ground truth for therefore only a few leaves from those plants can be compared.

Procedure

Locating the leaves in 3D on a grey surface model can be hard because the central section of the plant or dense leaves appear as lots of patches with only the shape to indicate if it is a leaf.

For each leaf on the ground truth model the corresponding leaf from the mapped model was located. The leaf ID was counted as the correct leaf and anything else was considered incorrect (this includes the faces where no hyperspectral data existed in the first place).

Each Mapped leaf ID plant has between 10 and 40 leaves and as long as there was ground truth for the leaf then it was included in the evaluation. The evaluation is an accuracy score from 0 to 1, where 0 is completely incorrect and 1 is a perfect match between the ground truth and the mapped data for each leaf. The data is mapped to patches; where each patch contains a group of triangles that collectively forms the plant model. The number of patches that match are counted and then divided by the total number of patches representing the leaf. Table (5.1) shows some selected leaf data of the plants as an illustration. Two plants from each variety were selected and then individual leaves were selected to represent the minimum, maximum and mean accuracy leaves. This allows for an overview of the data from each variety across a drought and control set from day 3 (Table 5.1). For the BSP14 plants the mean leaf accuracies range from 73% to 80% while there is a slightly wider range for the mean leaf accuracies Sonata plants (61% to 86%), However the range of the leaf accuracies for both varieties are similar (48% to 65% spread).

Table 5.1: Statistics of accuracy scores (0-1) of correctly mapped patches per leaf compared to the ground truth from four selected plants. The table displays the differences between Sonata and BSP14. Top two rows are BSP14 plants and bottom two rows are Sonata plants. Plants 4 and 13 are drought plants, plants 12 and 16 are control plants. This is a small subsection of plants for an overview of the statistics between the two varieties.

Plant	Min leaf ratio	Max leaf ratio	Mean leaf ratio
4 (BSP)	0.43	0.95	0.73
12 (BSP)	0.52	1.00	0.80
13 (Sonata)	0.35	1.00	0.61
16 (Sonata)	0.48	1.00	0.86

Boxplots have been included to display hyperspectral data mapping accuracy for all of the leaf data across two varieties. Figure 5.31 is the boxplot for 6 selected

BSP plants, 3 drought and 3 control plants. Then Figure 5.32 displays all of the plants for BSP14 and Sonata (except BSP14 plants 3 and 9). Due to the ground truth selection the Sonata plants have a reduced number of leaves that were compared.

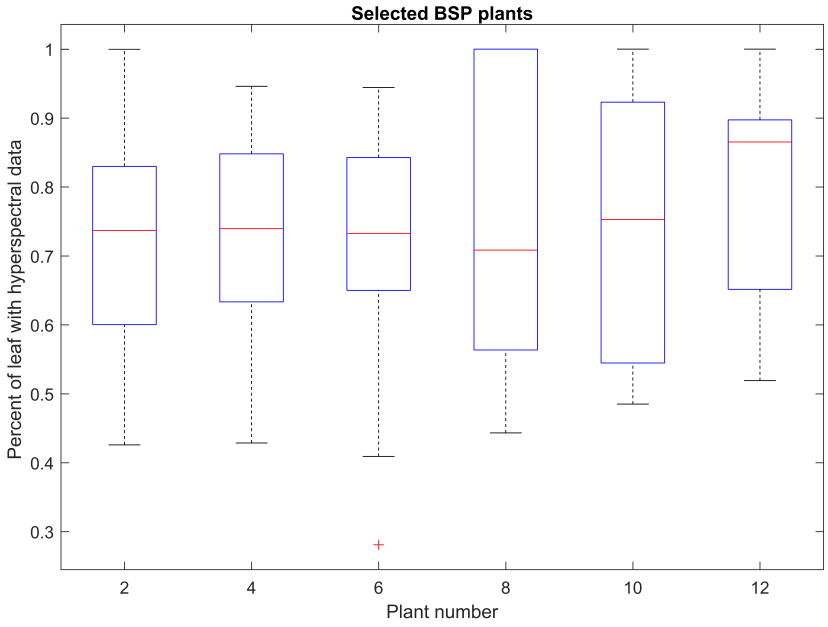

Figure 5.31: Six plants have been selected (3 drought (2,4,6) and 3 control plants(8,10,12)) to display the correctly mapped hyperspectral data and the variation between plants of the same variety. Each plant has a different amount of leaves; boxplot displays the summary statistics.

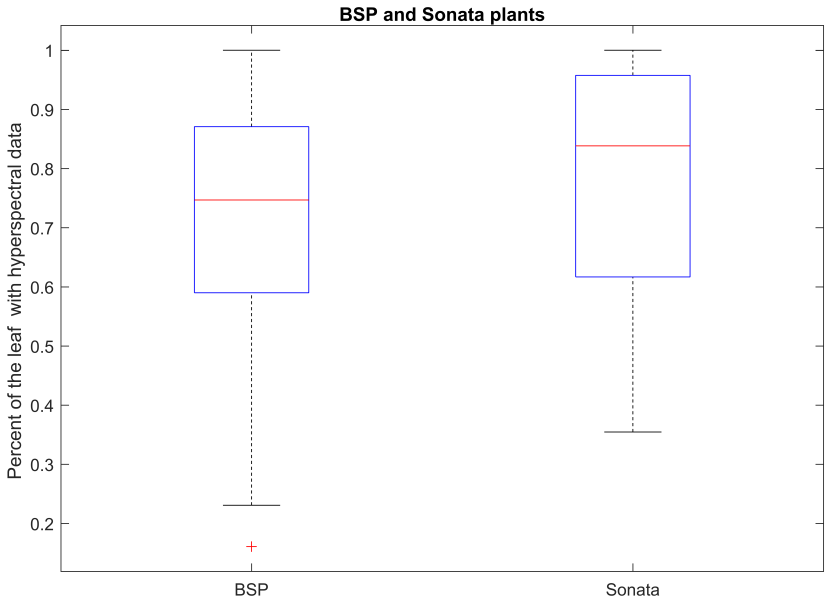

Figure 5.32: The leaf percentage accuracy across all plants of two varieties has been measured. BSP and Sonata are compared. There are 166 BSP14 leaves and 24 Sonata leaves.

5.9 Conclusions

This Chapter developed an approach to mapping 2D hyperspectral data onto a 3D plant model. This has been done to supplement the hyperspectral data with extra information about the leaves in 3D, such as the height, distance from the centre of the plants and the leaf orientation angles. This could be used to improve the hyperspectral data quality, and will be examined more in Chapter 6.

During the process of mapping the data from 2D to 3D, the 2D ground truth labels were used in order to know where the hyperspectral data came from, and if it was mapped to the correct leaf in 3D (by comparing the 2D to the 3D ground truth). This does mean that the ground truth in 2D would be needed. However if this method is used on a live system then instead of the 2D ground truth the level set segmentation for the leaves could be used from Chapter 4. This was not used here in order to isolate testing from earlier sections of the pipeline.

Table 5.1 and boxplots (5.31, 5.32) from the evaluation display statistics about the accuracy of the mapped data. The accuracy represents the successful map-

ping of hyperspectral data on faces within the leaf area (selected from the 3D ground truth). The extra faces that do not have the leaf data is labelled pink for visualisation. This is the areas where the hyperspectral data did not exist for the 3D model, where the hyperspectral image was cut off (cropped) or there is misregistration.

Table 5.1 summaries statistics for two BSP14 and two Sonata plants, the means vary between the varieties - BSP14 have means (0.73-0.80) while the Sonata plants are (0.61 - 0.86) which suggests the leaf IDs were mapped more consistantly for the BSP14 plants, however the boxplot will give a better result of the statistics because this includes more data.

The boxplots display the difference between the plants and between two varieties. Figure 5.32 displays the BSP14 leaf accuracies and the Sonata leaf accuracies. The BSP14 boxplot has a lower overall median accuracy compared to Sonata however BSP14's inter-quartile range is smaller, which means the leaf mapping accuracies are slightly more compact than Sonatas. The whiskers for the BSP14 is considerably longer than Sonata which indicates a larger distribution of leaf mapping accuracies.

The 3D information can be used to select good quality leaves for further analysis based on certain criteria. The leaf angle has been shown to affect the spectral signature therefore selecting a leaf parallel to the camera would potentially be good. The plant has better lighting near the top centre compared to lower down due to the position of the lighting and the camera position. Also mentioned previously shadows affect the spectral signature. 3D information now mapped to the hyperspectral data will allow for intelligent data sampling for use in further analysis: this will now be considered in Chapter 6.

Chapter 6

Using 3D information to select leaves from which to measure hyperspectral data

Chapter 5 involved 3D reconstruction of the strawberry plants and mapping hyperspectral data onto the 3D models. This then allows extra 3D information to be retrieved and used to select leaves from which to take hyperspectral measurements.

This Chapter looks at selecting high quality leaves for analysis. This would be leaves near the top of the plant that are clearly imaged by the camera, and the top leaves would also be well lit by the light source. It would also be beneficial if the leaves are as flat as possible and facing towards the camera.

There are many potential measures that could be used to select 'high quality' leaves, some potential candidates are listed below.

1. Leaf angle
2. Leaf height
3. Leaf distance from centre of the image
4. Size of the leaf
5. Leaf flatness - inclination from horizontal position across the leaf surface
6. Brightness of reflectance data

This Chapter will consider such options and develop a set of suitable measures which help identify a 'good quality' from which to take a hyperspectral measure.

6.1 Criteria for selecting leaves from the 3D model

There are many potential measurements that can be computed to help select the 'best quality' leaves.

1. Leaf angles would help to determine how the leaves are orientated relative to the camera; for hyperspectral images the angle can influence the hyperspectral profile.
2. The leaf height is useful to select leaves near the top of the plant and disregard the lower leaves that will be more likely to be in a shadow, or dark, or occluded.
3. The leaf distance from the centre of the image could be helpful; the lights are directed towards the middle of the plant and large plants might suffer from poor illumination on the outside leaves.
4. The size of the leaf would not in fact affect the best leaves selected because there would still be a lot of hyperspectral information on a small or large leaf; except in extreme cases of eg. very young leaves.
5. Leaf flatness looks at the leaf inclination from the horizontal plane across it's surface. Each leaf is comprised of a number of patches and each patch will have an angle. The leaf flatness would be good to know the 3D surface texture.
6. Brightness of the reflectance data will decrease with leaves lower on the plant; if the leaf height is taken into account then the leaves selected should already account for this.

From the list there are three factors that would be beneficial for selecting optimum data: Leaf angle properties (inclination and flatness) (A), the leaf height (B) and the leaf distance from the centre of the plant (C).

The leaf angles can influence the hyperspectral data by increasing the red wavelengths and decreasing the green and near-infrared wavelengths (Chapter 5.2), therefore it would be useful to know the angles and potentially select the flattest leaves (5). The height of the leaves from the top of the plant will affect the lighting on the leaf, via potential shadows from leaves above, and also the leaves near the bottom may not have much hyperspectral data mapped onto the leaf due to the top down view of the hyperspectral images ie, they are occluded. The leaves near the centre of the plant will have more consistant lighting due to the lighting

arrangement on the imaging system.

Figure 6.1 displays the flowchart for Chapter 6, to explain the process.

Figure 6.1: Flow diagram for selecting leaves in Chapter 6.

6.1.1 A: Finding leaf angle and flatness

The first challenge is to determine the leaf orientation (1) and the leaf flatness (2). To do this using the 3D data from Chapter 5 the approach must:

1) Select all the vertices from the leaf, fitting a plane to the points and comparing this plane with the imaging XZ plane in order to find the general inclination of the leaf compared to the surface/table plane.

2) Find the angles for *all* the patches in the leaf and compare the angles to quantify the 'flatness' of the leaf. For example if most of the patches have a similar angle then the leaf is considered flat, however if there is a variation in the angles then this indicates there is a curve in the leaf.

6.1.1.1 Development of overall leaf angle quantification

An angle can be found between two planes by finding the normal vectors of both planes. The normal vectors are tangent to the planes; therefore, the angle between the two normals is the angle that can be used to quantify orientation, and if one plane is the horizontal image plane, it can measure inclination of the leaf. In this case the leaf and the table surface (horizontal/XZ plane) are the two planes. The leaf plane will be an approximation because the leaf might be curved or a few patches could be misaligned. The leaf flatness will also be explained in Section 6.1.1.2 to find the leaf structure.

Least squares is used to find a plane of best fit for the vertices. The mean of the sample points is found, assuming the mean point lies on the plane. Then the difference between the mean and the vertex points are selected and the covariance matrix is computed. The eigenvector with the largest eigenvalue is in the direction of the largest variation (which should be along the plane); therefore the other eigenvectors are orthogonal. The normal is found from selecting the eigenvector with the smallest eigenvalue.

The imaging plane for the camera in this case is the XZ plane; therefore the normal for this plane is (0,1,0). Figure 6.4 displays the normal vectors with the leaf and horizontal plane.

The last step is to find the angle. The dot product and cross product are calculated, the normals need to be unit length, then arctan can be used to compute the angle (in radians - which is then converted to degrees). The angle will be in the

range 0 - 180, but only 0 - 90 degrees is required to the horizontal, because the leaf angle of interest is between the horizontal and vertical planes (90 degrees); therefore if the angle is more than 90, subtract the angle from 180. By way of example, the calculated angle for BSP14 plant 12 leaf 15 is 27.8 degrees; see Figures 6.2 and 6.3 for the displayed leaf and leaf vertices, and Table 6.1 for the leaf angle.

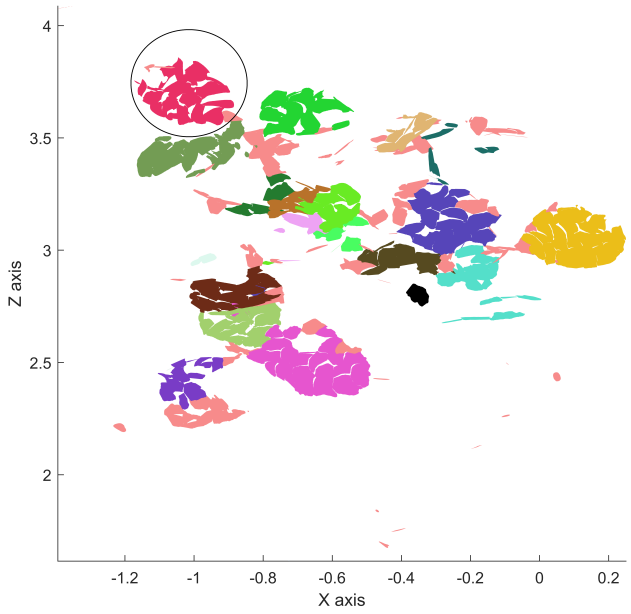

(a) Plant 12 top down view

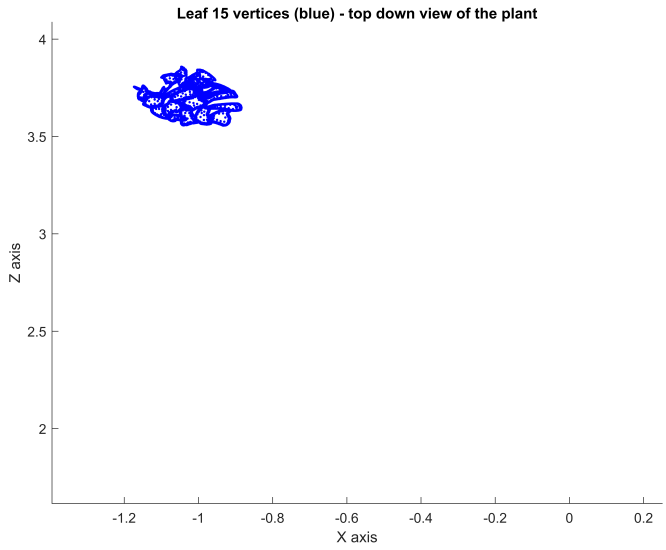

(b) Leaf 15 vertices top down view.

Figure 6.2: (a) BSP14 Plant 12 with a circle around leaf 15. (b) leaf 15 vertices from a top down view along the X axis and Z axis.

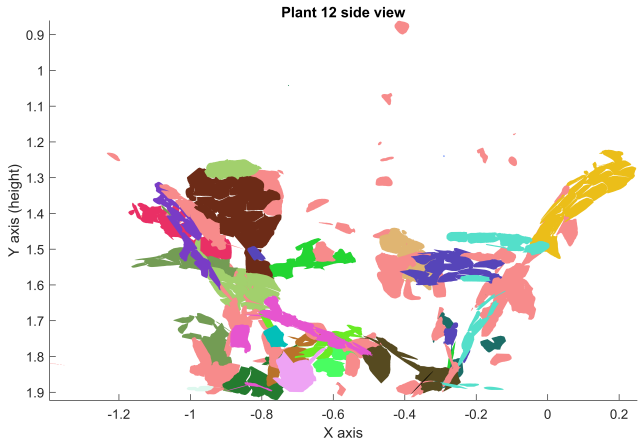

(a) Plant 12 side view

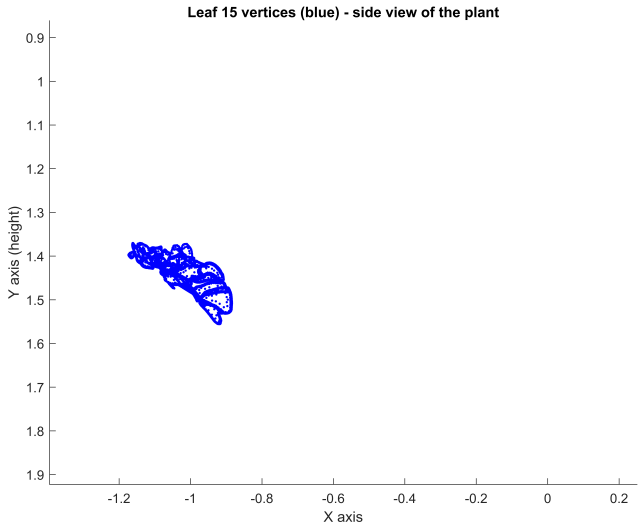

(b) Leaf 15 vertices side view.

Figure 6.3: (a) BSP14 Plant 12 and (b) leaf 15 vertices from a side view where the Y axis represents the plant height.

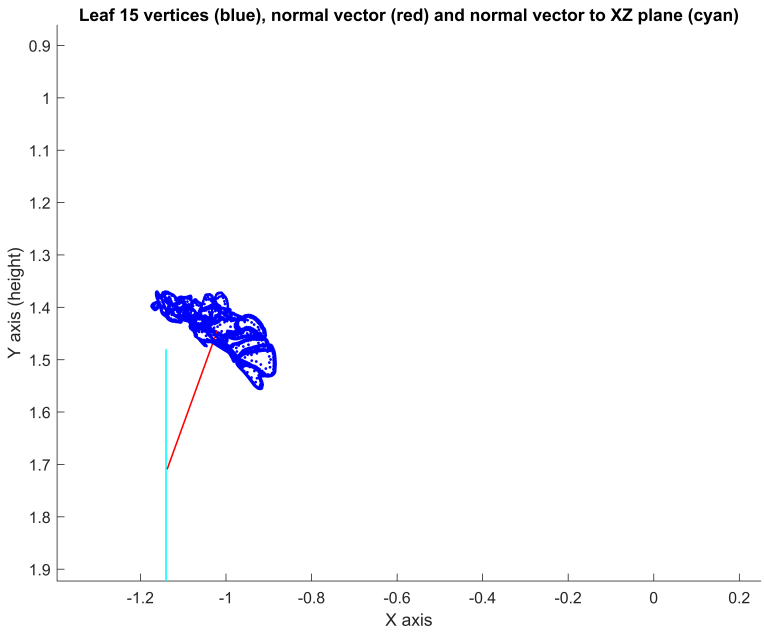

Figure 6.4: Leaf 15 vertices (blue) with the normal vector (red) and normal to the XZ plane (cyan)

This process computed the leaves on one particular plant for an example. The leaves in Table 6.1 have the corresponding leaf angle computed through the approach as explained. These are found by fitting a plane to the vertices for the leaf angle. The following Section considers leaf shape through the patches that create the leaf, in order to identify how flat the leaf is.

6.1.1.2 Leaf patch angle quantification to determine the leaf 'flatness'

The leaf data is comprised of a group of patches. The individual patches are flat, and together they create the appearance of the leaves. The leaves can be flat, curved or curling. Also there might potentially be a random noisy patch at the wrong angle compared to the rest of the patches due to a problem in the process when the patch was created.

A flatness score would therefore require finding the individual patch angles and comparing the list of angles to see if they are similar or not. The patches could have a similar angle (therefore, flat), a few distinct angles (wavy or originating from a technical problem) or lots of smoothly varying angles over a potentially wide range(curved). The angles for all the patches are needed, and then the angles can be analysed to see if the leaf is flat or curved.

Table 6.1: Plant 12 leaf number and the leaf angle. The angles are ordered from smallest to largest, where 0° is parallel to the camera plane and 90° is tangential to the camera.

Leaf number	Angle
6	16.8
14	22.8
15	27.8
1	29.5
9	38.0
8	38.0
11	44.0
19	54.7
13	61.3
2	64.0
16	65.8
10	70.6

First, to find the patch angles the same process as for flatness is used to find the angles for the leaf only this time the angles are calculated for the individual patches in the leaf. Then a novel algorithm to quantify flatness is presented below.

As an example, leaf 15 will be used again. The angles for the patches in leaf 15 are listed in Table 6.2. To note, while this relates to leaf 15 in Table 6.1 the overall leaf angle was calculated from fitting a plane to the vertices in the leaf (Section 6.1.1.1), therefore it can vary from the patch angles in Table 6.2. This can vary because the plane is fitted to the leaf vertices of the leaf (point cloud), however with the patches there can be lots of local variation. The overall leaf angle for leaf 15 is 27.8° however the mean angle for all the patches in leaf 15 is 42.5°.

In general if the patches were aligned with no noise, no local variation or missing patches then the overall leaf angle should be very similar to the average leaf patch. Table 6.3 is included to illustrate the difference between the leaf angle and average patch angle. Seven of the leaves have a similar angle (with a 10 degree tolerance). Five of the leaves are above that level. Even though they are considered different, the largest difference is 30 degrees. Therefore both angles are needed because they are not always the same.

If patches are missing and the leaf is curved or folded then the overall mean patch angle would be different from the leaf angle plane. For example Figure 6.5 is a diagram with two curves (green) and patches (blue). The overall angle for both curves is the same, but the local average patch angle is different. The top Figure

Table 6.2: Leaf 15 patch angles

Patch number	Angle (degrees)
1	71.1
2	53.1
3	58.3
4	73.9
5	30.5
6	33.8
7	19.6
8	8.0
9	67.7
10	18.7
11	64.2
12	35.8
13	43.8
14	38.5
15	22.8
16	38.0
17	45.3
18	11.1
19	58.2
20	58.2
21	41.3

Table 6.3: Comparison of the overall leaf angles and average leaf patches from individual leaf patches. Ordered by the difference, from the smallest to largest.

Leaf	Patch average	Leaf average	Difference
11	44.3	44.0	0.3
19	55.3	54.7	0.6
13	60.5	61.3	0.8
8	37.1	38	0.9
1	28.0	29.5	1.5
9	31.3	38	6.7
6	25.9	16.8	9.1
14	34.0	22.8	11.2
2	52.0	64	12.0
15	42.5	27.8	14.7
16	45.2	65.8	20.6
10	40.7	70.6	29.9

includes patches that are aligned and matching where the curve is. The bottom Figure displays some patches missing and this would influence the average patch angle. The figure is shown to explain how the two angle metrics can be different, strawberry leaves are usually more flat than the curves in the diagram however patches might still be missing due to occlusion and noise during the reconstruction as well as the input images for reconstruction being quite limited in both number and viewing angle.

Both measures can still be used to identify different features.

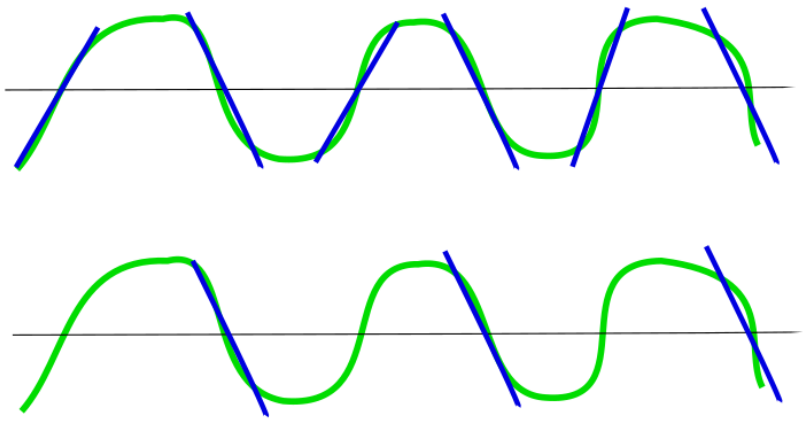

Figure 6.5: Diagram of two curves (green) and patches (blue) to display how a leaf can have the same overall angle but the average patch angle can be different. The top Figure is when the leaf and average patch angle would be the same or very similar. The bottom Figure has missing patches, and this will influence the average patch angle to be larger than the leaf angle. The black line represents the leaf overall angle.

Figure 6.6 shows a leaf and the patch angles, distribution of angles and the KDE graph of the peaks. Figure 6.6a displays a RGB image of leaf 15 on plant 12. Figures 6.6b and 6.6c are the reconstructed leaf patches in Meshlab with the front view, similar to 6.6a, and the side view to show the variation in the patch orientations.

The histogram (Figure 6.6d) displays the distribution of patch angles, grouped in degree increments. There might be noisy data or anomalies; therefore any bins with a frequency of 1 are removed from the data. There are usually 10-30 patches per leaf therefore removing a few patches should still leave plenty which are suitable

to calculate flatness. The next histogram (Figure 6.6e) reflects the new data once these potential anomalies have been removed. The kernel density estimation (KDE) (Figure 6.6f) estimates the underlying probability density function of the data [116]. The samples are each smoothed using a Gaussian kernel and summed for a total estimate. There is also a bandwidth to control the smoothness of the data, the smoothing can artificially increase the variance and lower the peak. The peaks are where the majority of the patch angles are; if a leaf is flat then all of the patches should have similar angles. However, if the leaf is a smooth curve then the angles would be distributed over a wide range, this would still represent one peak but the spread of the data would be large.

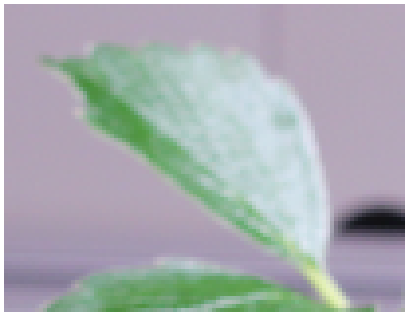

(a) Plant 12 leaf 15 rgb image

(b) Plant 12 leaf 15 meshlab image

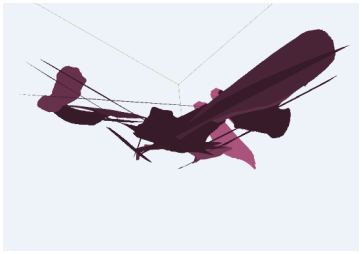

(c) Plant 12 leaf 15 meshlab image side view

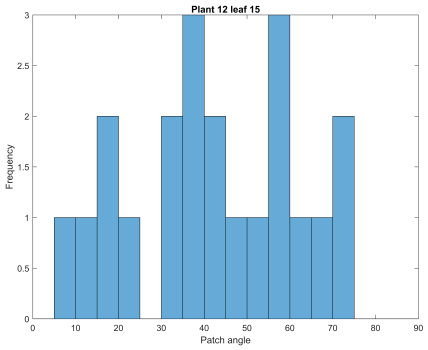

(d) Plant 12 leaf 15 histogram of the angle data

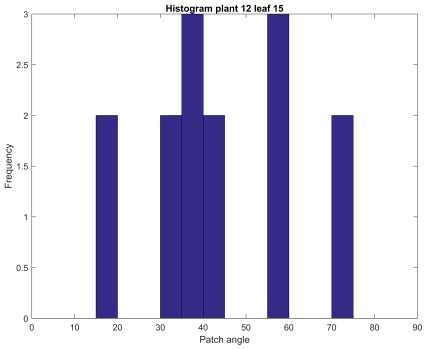

(e) A histogram of the edited data to remove angle bins with 1 frequency.

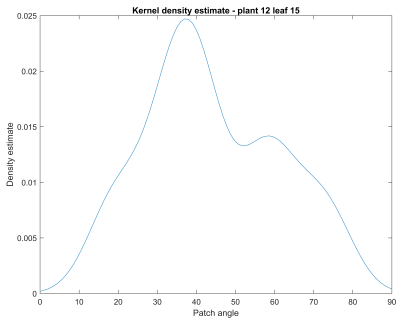

(f) Kernel density estimate of the angles to display the peaks.

Figure 6.6: Plant 12, leaf 15 displaying the leaf in RGB (a), the reconstructed leaf front view (b) and side view displaying a fold in the leaf (c). The histograms (d,e) and KDE graph (f) show the angles, the range and the frequency of the patches in the leaf.

The peaks on the KDE profile are counted by finding any local maximum points (where the two point beside the maxima are both a lower value). If there are three or more peaks then the leaf is determined to be curved (or not flat). If there are two peaks and the range of the angles are less than 25 degrees then the leaf is considered sufficiently flat, otherwise the leaf is considered not flat. Finally if there is only one peak the range is still checked - if the range of angles is less than 35 degrees then the leaf is flat otherwise it is not flat. For example, the leaf can have one peak but if the angles vary by 50 degrees then the leaf could be curved.

If there are one or two peaks in the KDE then the leaf could still be flat enough, this is why up to three peaks are selected. Also 35 and 25 degrees are empirically selected as an estimate for when the leaf would be considered not flat.

The peaks are counted from the KDE, and the range of angles is selected from the histogram. The KDE is smoothed, which means the spread can appear larger than it is. In general the more peaks there are the less likely the leaf is to be flat, also the bigger the spread of angles in the histogram is also an indication that the leaf is not flat.

So the rules for flatness are as follows:

- Number of peaks: 1) if angle range <35 then leaf = flat
  otherwise leaf = not flat
- Number of peaks: 2) if angle range <25 then leaf = flat
  otherwise leaf = not flat
- Number of peaks: 3+) leaf = not flat

Three examples with one (Figure 6.7) , two (Figure 6.6) and three peaks (Figure 6.8) are displayed. The approach results in two new properties per leaf: a flatness label, and a count of KDE peaks.

The side view profile for leaf 9 (Figure 6.7c) and the histogram (Figure 6.7e) visually shows a flat leaf; there is one peak and the range of angles are relatively close together. Leaf 15 (Figure 6.6) displays a folded leaf and there are two peaks, this would be considered not flat. Then leaf 16 (Figure 6.8) is a noisy leaf, with a wide range of angles. For this data it would be hard to assign a variable with a continuous value to measure how flat the leaf is because the leaf patches are noisy. This could be due to the limited number of images for the reconstruction process.

A binary value of 1 or 0 will be used because for some leaves the patches do not align and measuring this variation in a curvature value could determine a

higher value than the leaf actually is. If the data was collected again and more images were captured then the leaves would have less variation/noise. In this case the Gaussian or mean curvature could be measured. An advantage for using the binary score in this instance is to determine a general leaf shape rather than a fine grained description.

Table 6.4 displays the leaf number, whether the leaf is considered flat or not and how many peaks there are of similar angles.

(a) Plant 12 leaf 9 rgb image

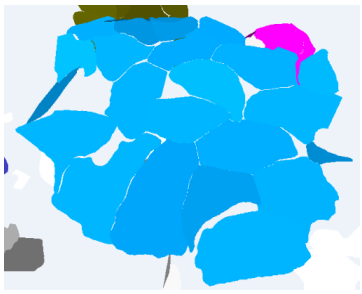

(b) Plant 12 leaf 9 meshlab image

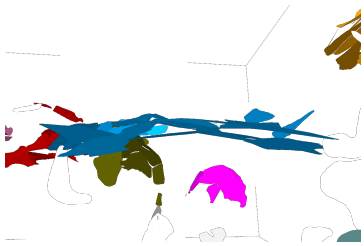

(c) Plant 12 leaf 9 meshlab image side view

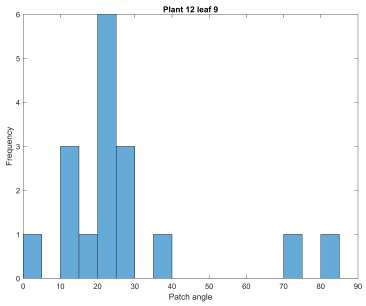

(d) Plant 12 leaf 9 histogram of the angle data

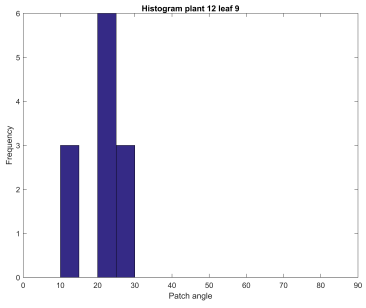

(e) A histogram of the edited data to remove angle bins with 1 frequency.

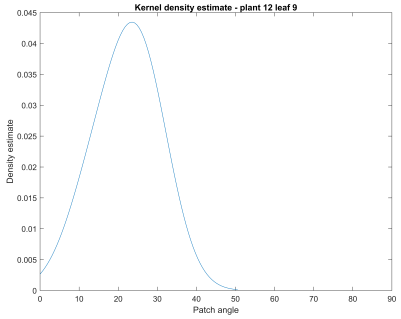

(f) Kernel density estimate of the angles to display the peaks.

Figure 6.7: Plant 12, leaf 9. The histograms and KDE graph displays the angles, the range and the frequency. The leaf is smoothly curving therefore there is only one peak.

(a) Plant 12 leaf 16 rgb image

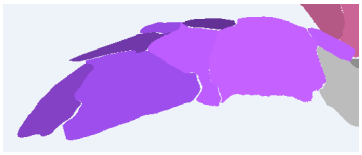

(b) Plant 12 leaf 16 meshlab image

(c) Plant 12 leaf 16 meshlab image side view

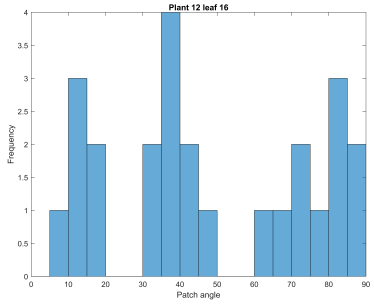

(d) Plant 12 leaf 16 histogram of the angle data

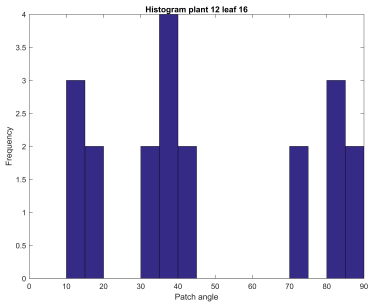

(e) A histogram of the edited data to remove angle bins with 1 frequency.

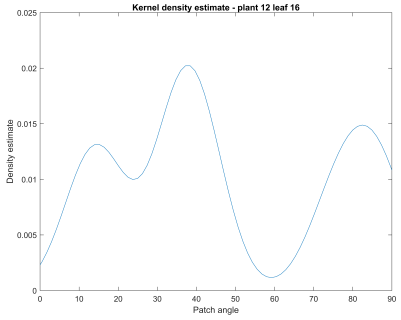

(f) Kernel density estimate of the angles to display the peaks.

Figure 6.8: The histograms and KDE graph displays the angles, the range and the frequency

Table 6.4: Plant 12 - leaf flatness and peaks. Ordered by the leaf flatness score

Leaf Number	Flat	Peaks
1	1	1
8	1	1
9	1	1
13	1	1
19	1	1
6	1	2
2	0	1
11	0	1
14	0	2
15	0	2
10	0	3
16	0	3

6.1.2 B: Height of the leaf

The leaf height on the plant determines how close the leaf is to the camera and the top of the plant. It has been stated before that the leaves near the top of the plant have been clearly imaged, they should not be in a shadow or covered by an overlapping leaf, and they would therefore be sensible leaves to select.

Due to the registration in Chapter 5, the height of the leaves aligns along the Y axis of the 3D data. Therefore, the leaf vertex Y-coordinate is all that is needed. The average of the Y-coordinates for the leaf is calculated, and this is used to represent the leaf height.

Figures 6.9 and 6.10 visualise the leaf heights. The values are stored in a Table for each leaf and this will be used as a measure for determining the best leaves to be selected. Table 6.5 lists the leaf numbers and leaf height.

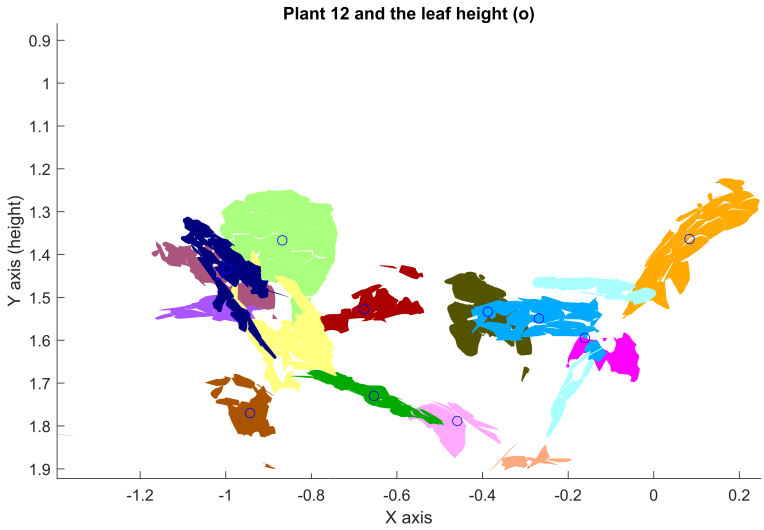

Figure 6.9: BSP14 Plant 12 and the leaf height measurement (blue circles).

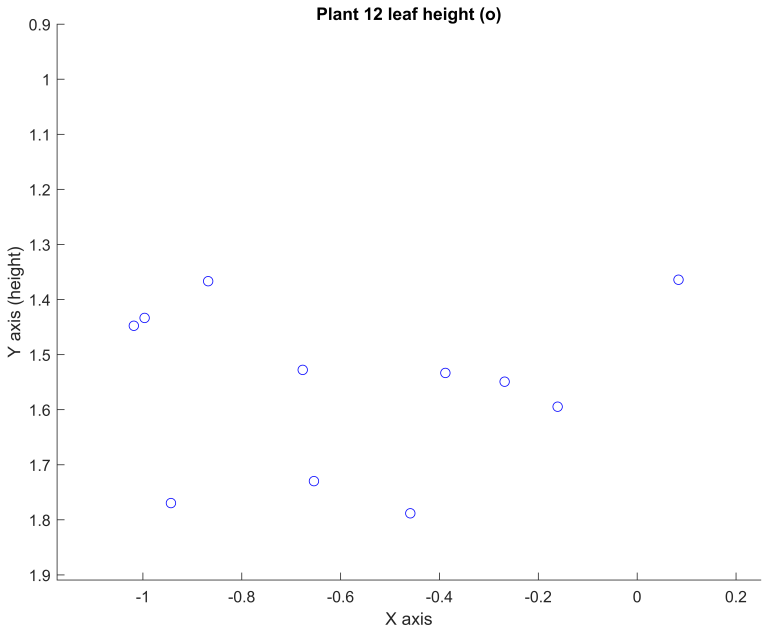

Figure 6.10: BSP14 Plant 12 leaf points without the plant to see the points used to measure height.

Table 6.5: Plant 12 leaf number and the leaf height. Lower values are closer to the top of the plant. Ordered by the height from the top of the plant.

Leaf number	Height
11	1.364
2	1.367
19	1.433
15	1.448
14	1.528
13	1.533
9	1.549
10	1.595
6	1.718
1	1.730
16	1.770
8	1.788

6.1.3 C: Distance of the leaf centre to the plant centre

The plants are centred under the camera and the lights are pointing down towards the plant, they are angled to the middle of the camera system set up area. If the plants are quite small (eg. BSP - plants 1,...12) then most or all of the plant should be well lit; however, with large plants (Sonata - Plants 13,...18) the edge of the plant and leaves could have low lighting due to the hardware set-up and the positioning of the lights. Therefore the distance of the leaf from the centre of the plant, and hence imaging system, could be considered when selecting the 'best leaves'.

In this case the centre of the plant and the centre of the imaging platform are the same since the plant is positioned in the centre of the image platform. However if the plant is positioned off centre then the platform centre would need to be the focus due to the location of the lights.

The centre of the plant and the centre of each leaf is calculated, then the distance between the plant and the leaves is measured. There are a few ways of selecting the centre point from a group of vertices. It is possible to use the mean of the points, the median of the points, or weighted versions of these averages.

The centre points calculated from Chapter 4 used the mean of the points and this is used again here. For the set of plant vertices $S = (x_1, z_1), (x_2, z_2), ...(x_n, z_n)$ the average x and z values are found $(X_P, Z_P) = (\frac{1}{n}\sum x_i, \frac{1}{n}\sum z_i)$. For the set of

leaf vertices $T = (x_1, z_1), (x_2, z_2), ...(x_n, z_n)$ the average x and z values are found (X_L, Z_L). Then the distance is computed using Euclidean distance.

Figure 6.11 shows the plant and leaf center points for plant 12. These are used to measure the distance of the leaf to the center of the plant. This will be used as a further measure for selecting the best leaves.

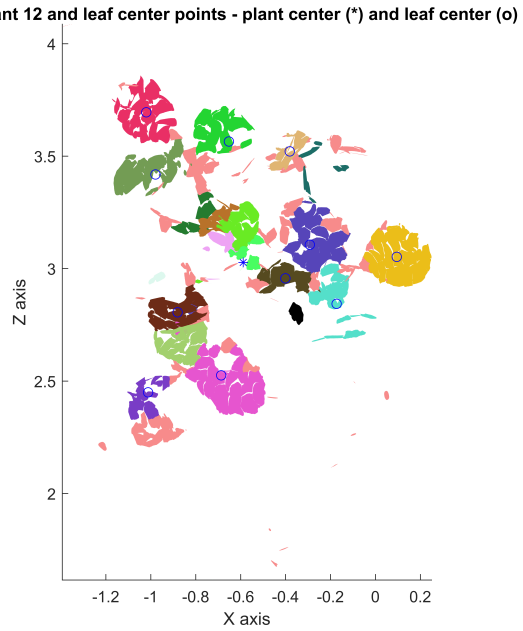

Figure 6.11: Plant 12 center point (*) and the leaf center points (o) have been displayed. A top down view is shown.

6.1.4 Selecting the best leaves to be analysed

After finding the leaf angles, heights and distance from the center of the plant the next step is to use this information in order to select leaves to use for subsequent hyperspectral measurement. Each measure A: angle, B: height and C: distance will be sorted in order and scored from 1 to N. Then A B and C will be weighted to decide which three leaves will be selected.

Table 6.7 shows a list of the leaves along with the angles, heights and distances for one plant. The angle, height and distance have been individually sorted from smallest to largest. The smallest angle is desired because it would be parallel with the XZ plane (surface). The smallest height in this case means the top of the plant

Table 6.6: Plant 12 - Leaf distance to centre. The distance is sorted in order from smallest to largest, where the smallest is the closest leaf to the centre.

Leaf	Distance
6	0.143
8	0.201
9	0.307
2	0.367
10	0.455
1	0.512
13	0.535
14	0.541
16	0.553
11	0.684
19	0.716
15	0.795

which is closest to the lights/camera. The smallest distance is the point closest to the center of the plant and hence, image.

Note the leaf flatness is not taken into consideration in the score equation. It is currently a binary score, and it would be a challenge to balance the equation to include the flat score while maintaining the leaf angle and height as important measurements. While a binary score is not an efficient way to measure the leaves, due to the data and the reconstruction some of the leaves are noisy, this can be seen in Figures 6.6,6.7 and 6.8.

Table 6.7 lists the angles, heights and distances for the leaves of the BSP14 plant 12 3D model. Then Table 6.8 lists the summed scores for the leaves. The flat score and number of peaks are also included for reference. Columns 2,3 and 4 are summed for a total score (column 5), and the Table is ordered by column 5.

To provide a quality measure, a score is computed from the 3 metrics. The score is initially calculated from A + B + C = S; smaller scores are better. This is a simple method for including all three measures; variations on this are possible and will be examined below. Table 6.8 lists the scores and the leaves sorted in the order of the smallest to largest score. If the best three leaves were selected, the system would select leaves 6, 14 and 9 in this case.

The equation could be weighted more towards particular measures. The selection process and weighting could change which leaves are selected. For example Table 6.9 only uses the angles and height to determine the final score. If the top three

Table 6.7: Plant 12 - Leaf angle, height and distance to centre with the corresponding score. The angle, height and distance is sorted in order from smallest to largest. In this Table the three Sections are sorted independently for each measure; This shows that scores for each leaf can change between measures.

Leaf	Angle	Leaf	Height	Leaf	Distance	Score
6	16.8	11	1.364	6	0.143	1
14	22.8	2	1.367	8	0.201	2
15	27.8	19	1.433	9	0.307	3
1	29.5	15	1.448	2	0.367	4
9	38.0	14	1.528	10	0.455	5
8	38.0	13	1.533	1	0.512	6
11	44.0	9	1.549	13	0.535	7
19	54.7	10	1.595	14	0.541	8
13	61.3	6	1.718	16	0.553	9
2	64.0	1	1.730	11	0.684	10
16	65.8	16	1.770	19	0.716	11
10	70.6	8	1.788	15	0.795	12

Table 6.8: A list of leaves with the corresponding angles, heights and distances scores for plant 12, including a summed total score.

Leaf	Angle(A)	Height(B)	Distance(C)	Total score (A+B+C)
6	1	9	1	11
14	2	5	8	15
9	5	7	3	15
2	10	2	4	16
11	7	1	10	18
15	3	4	12	19
8	6	12	2	20
1	4	10	6	20
13	9	6	7	22
19	8	3	11	22
10	12	8	5	25
16	11	11	9	31

leaves are again selected then this would change from numbers 6, 14 and 9 to 15, 14 and 11. Only one leaf remains in the same top three across the two Tables; clearly choice of metrics is important. Table 6.9 shows an example of weighting the measurements in an extreme case, where A and B are weighted 1 and C is weighted 0.

Some measurements might have more influence on what determines a good leaf

Table 6.9: Illustration of the effect of using Angle and Height parameters only in the total score. A list of leaves with the corresponding angles, heights and distances scores for plant 12. The final score only includes A+B (Angle and height)

Leaf	Angle (A)	Height(B)	Distance(C)	Total score (A+B)
15	3	4	12	7
14	2	5	8	7
11	7	1	10	8
6	1	9	1	10
19	8	3	11	11
2	10	2	4	12
9	5	7	3	12
1	4	10	6	14
13	9	6	7	15
8	6	12	2	18
10	12	8	5	20
16	11	11	9	22

than others, and a weighted system could be helpful. The leaf angle and height are known to affect the hyperspectral data; however the distance from the centre of the plant could be less of an influence for the data. It has been demonstrated (Table 6.9) that an extreme weighted score will change the best three leaves; however, a small change in the weighted score system could be appropriate, distance is still a good measure to include but perhaps with less weight than the other parameters. Therefore, we shall explore what happens when the scores are weighted with a small difference in the weights.

The parameters should be tested to see how much influence they would have over the selection of the top three leaves if the weights were increased or decreased. One sensible option is to change the weights of one parameter at the time and keep the remaining two constant ie at 1 (a balanced weight).

The parameters have been tested by changing one weight at a time:

$$S = \sum(\alpha A, \beta B, \gamma C) \tag{6.1}$$

Where A is the angle, B is the height and C is the distance. For example, one weight (α) varies from 0 to 2 with 0.1 increments and two weights ($\beta\gamma$) are held constant. i.e. $\alpha = 0 : 2, \beta = 1, \gamma = 1$

The graphs in Figure 6.12 display three example plants with the results of the

angle, height and distance influencing the change in the top three leaves that are selected. The default parameter set is $\alpha = \beta = \gamma = 1$ and for each weight the new selection of leaves are compared to the default selection of the top three leaves. For each graph the Y axis is numbered 0 to 3 where 3 represents the same top three leaves selected as the balanced equation and 0 is that none of the original three leaves are the same. The X axis is the weight value of the parameter under test which varies from 0 to 2 in 0.1 increments.

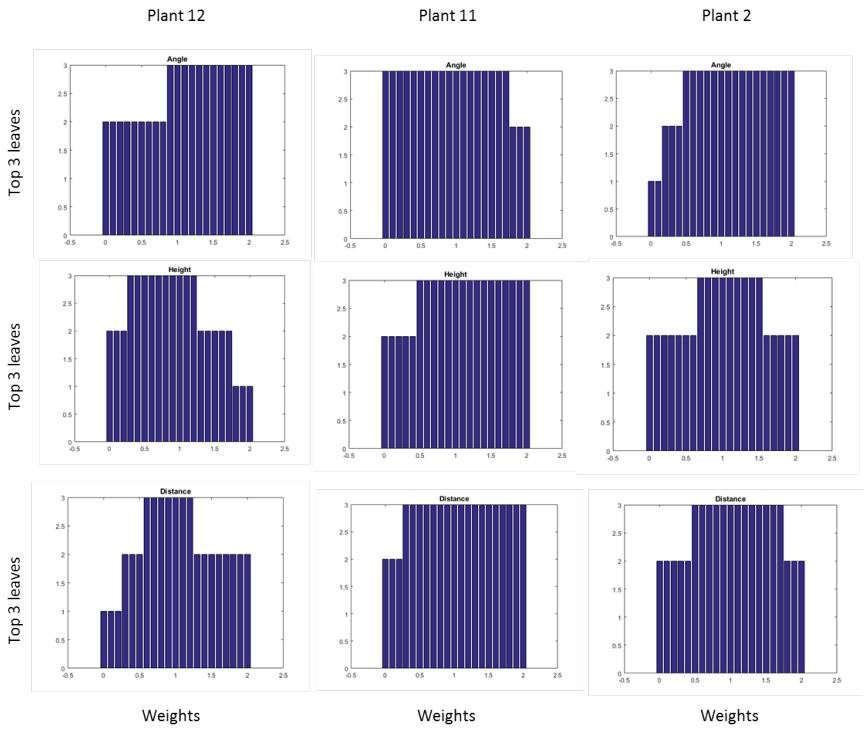

Figure 6.12: Graphs to display the number of top three leaves which remain the same when the weights are changed for the angle, height or distance parameters. Top row are the angles, middle row are the heights, and bottom row are the distance for the three plants.

The weights affect which top three leaves are selected, however the same three leaves remained the same for three plants when the weights varied slightly above or below 1. This means the weights can change the outcome by being increased or decreased by large amounts; however, values around 1 do not affect the chosen leaves much.

The lowest final score is the best quality leaf and the highest score is the worst option. The equation could be weighted (0.8,0.8,1) and this would mean that one of the measurements influences the final score less than the other two. The

weights can be increased, decreased or balanced. From the graphs above this would generally not affect the top three leaves.

The equation could also include the weights after a certain limit. Once the leaf angle is above an extreme limit the weight can be increased for the angle score, if the leaf height goes above a certain value then the height score could also be increased.

Equation 7.1 starts with the weights $\alpha = 0.8, \beta = 0.8, \gamma = 1$

.

Additional penalty for extreme values

This Section then adds an additional constraint so that once the leaf angle and or leaf height reaches a certain limit then they are penalised. This just increases the score and reduces the likelihood that the leaf would be selected. It effectively filters out extreme values entirely.

if A >50 degrees then $\alpha = 1.2$
if B >50 % then $\beta = 1.2$

Equation 7.1 is then re-weighted using the formula above. This would produce Table 6.10 which includes the ordered scores for the leaves. The top 3 leaves remains the same as the balanced Table 6.8 in a different order. This makes sense, as top three leaves should not contain many extreme values anyway. While the same top three leaves would be selected for the balanced equation and weighted equation in this case it might change slightly for different plants.

Following development and application of this score, the top three leaves selected are displayed in Figure 6.13. Leaves 6 and 9 also appear to be flat from the criteria used to define the flatness. This metric can be calculated automatically for all plant models in the experiments.

Table 6.10: Plant 12 - the weighted scores for the leaves and corresponding angles, heights and distances.

Leaf	Angle (A)	Height(B)	Distance (C)	Total score (weighted)
6	1	9	1	12.6
9	5	7	3	12.6
14	2	5	8	13.6
11	7	1	10	16.4
2	10	2	4	17.6
15	3	4	12	17.6
8	6	12	2	21.2
1	4	10	6	21.2
13	9	6	7	22.6
19	8	3	11	23
10	12	8	5	25.8
16	11	11	9	35.4

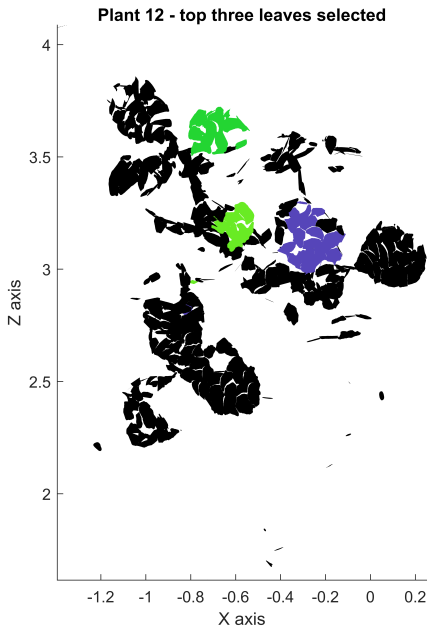

Figure 6.13: BSP14 Plant 12 - the top three leaves selected (6,9,14) from the weighted (and balanced) equations

The results are shown for one plant in Table 6.10, and the top 3 leaves have been displayed in Figure 6.13. Three leaves have been selected for each plant as a sample number of leaves to analysed. This could have been any reasonable number but a small number of leaves would still result in a lot of hyperspectral data that is selected, and 3 allows for some replication of samples.

6.1.5 Summary

The Chapter has discussed three leaf properties that can be measured from the 3D data.

A novel leaf scoring system was developed to automatically select leaves, from the automatic pipeline developed in Chapters 4-5. The flatness score was not included in the weighted scoring system, it would need to be balanced so the angle, height and distance are still regarded as important in determining the leaf order. The flatness could be incorporated by implementing a sensitivity analysis for the parameters to see how the different combinations influence the final leaf order.

The weighted scoring system includes a penalty section by increasing the parameter when the measurement reaches a limit and this affects the leaf score. This is one example approach, but it might not be the best option to penalise the leaves during the weighted system, instead the extreme values could be filtered out directly.

In the next Chapter the selected top leaves are analysed over the time series.

Chapter 7

Hyperspectral data analysis of the selected leaves over time

Chapter 6 developed an approach to locate the best (for a given quality measure) three leaves per plant to be analysed. The leaves are selected by determining the orientation and location of each leaf and then scoring the leaves from best to worst. The top three leaves would ideally have a shallow angle and be centrally located near the top of the plant.

This Chapter focuses on the leaves over the time series to see if there are any differences over the time during drought stress until visible signs appear.

7.1 Comparing independant, automated selection of leaves vs tracked selection of leaves

The automated pipeline from Chapters 5 and 6 is used to select the best three leaves for all of the plants. Chapter 5 reconstructs the 3D model, maps the hyperspectral data onto the model, then into hyperpatches. Chapter 6 takes this model and finds three selected leaf properties; the properties are measured and scored to select the leaves for day 3. The remaining question is how to select leaves from the other time points.

Option 1: Select the best three leaves, per plant, then manually match for the other days in the rest of the time series to select the same leaves. This is matched by manually inspecting the 2D ground truth images to find the correct leaves and

then manually ensuring that the correct leaves have been selected at each time point.

Option 2: Use the automated process for all of the days in the time series independently.

This raises the question of if the three leaves selected will be the same throughtout the time series. The effect of automatic selectioni on which leaves are used will first be investigated.

Effect of automatic selection

Figure 7.1 displays plant 12 on day 0,3,4 and 5 with the top three leaves selected by the automatic selection process. Day 3 is the original day selected and used in Chapter 6 therefore the leaf numbers relate to day 3 ground truth to easily match the leaves over the time points.

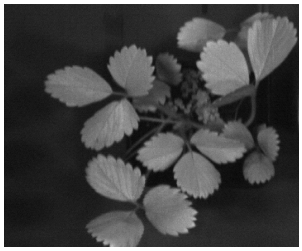

(a) Plant 12 day 0 top down view

(b) Plant 12 day 0. Top 3 leaves; 5 (green), 9 (blue), 14 (purple).

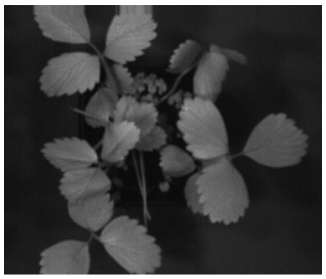

(c) Plant 12 day 3 top down view.

(d) Plant 12 day 3. Top 3 leaves; 6, 9, 14.

(e) Plant 12 day 3. Numbered leaves for reference.

(f) Plant 12 day 4 top down view.

(g) Plant 12 day 4. Top 3 leaves; 9 (purple), 14 (lime), 15 (green).

(h) Plant 12 day 5 top down view

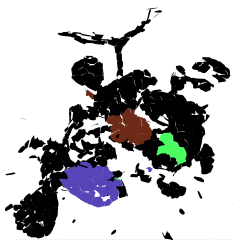

(i) Plant 12 day 5. Top 3 leaves; 1 (purple),7 (brown),10 (green).

Day	Number leaves
0	2
4	2
5	0

(j) Table of leaves that are the same as day 3.

Figure 7.1: Plant 12 time series. The left column is a top down view of the plant (days 0-4), the middle column is the 3D model with the top three leaves automatically selected. Days 0-4 include 2/3 same leaves selected. Leaf 6 is a young leaf in day 3. Day 5 has selected three different leaves, the plant has started to wilt.

The automated process selects the same two leaves from day 0 to day 4; leaves 9 and 14 (Figure 7.1). Leaf 6 in day 3 is a young leaf and it is tiny in day 0, also by the time leaf 6 is imaged on day 4 the leaf has moved slightly and is therefore not as good as leaf 15 using the automated selection.

Day 5 using the automated process selected three different leaves. This is also the time point where the first signs of drought appeared, which is likely to explain the change in selection.

The two approaches, fully automated vs tracked, can produce slightly different leaf selection however this is usually when the leaves have moved, have not been reconstructed in the 3D model or when visible signs of drought appear.

The leaves from the tracked method select the same leaves over the times series. This means the same leaves are analysed and this could include wilting leaves once they start to droop. Alternatively the automatic method selects slightly different leaves according to the scoring criteria in Chapter 6, because the leaves move slightly, or new leaves grow and therefore different leaves are now better. This means the healthier and better leaves will be selected and so it could potentially take longer to detect drought.

For this reason, the tracked-leaf method will now be used in this thesis to analyse the same leaves and maintain consistency.

The next Section looks at analysing data from the automatically selected, manually tracked leaves (from day three with the same leaves selected from day 0,4,5) and compares it with the initial data selected in Chapter 3.

7.2 Using automatically selected leaves to measure hyperspectral data by manual selection over time

Chapter 6 automatically finds the best three leaves for day 3. Then the best three leaves, per plant, are manually matched for the other days in the time series to select the same leaves. This is matched by overlaying the 2D ground truth images to find the leaves and then manually ensuring that the correct leaves have been selected at each time point.

The average of the mapped hyperspectral data for the leaves for each variety for

each day is calculated and plotted over the time series for the 16 plants (10 BSP, 6 Sonata). BSP14 plants 3 and 9 were discounted as they could not produce a 3D reconstructed plant.

The data is compared with the initial manually created plots in Chapter 3, presented again here for clarity in Figure 7.2a and 7.2b. For the initial manual analysis several leaves were selected per Sonata plant. The leaves appeared brightly lit and relatively uniform. Nothing else was known about the leaves at that point. The control plants spectral signatures revealed no obvious change from day 0 to day 5, whereas the drought plants have a slight shift in the signatures specifically visible in the near infra-red range (750 - 1000nm). Although quite small and so potentially insignificant, drought plants are expected to decrease in reflectance in the near infra-red (and slightly in the green wavelengths) over time. So this consistant change which is observed, although small, does fit with expectations.

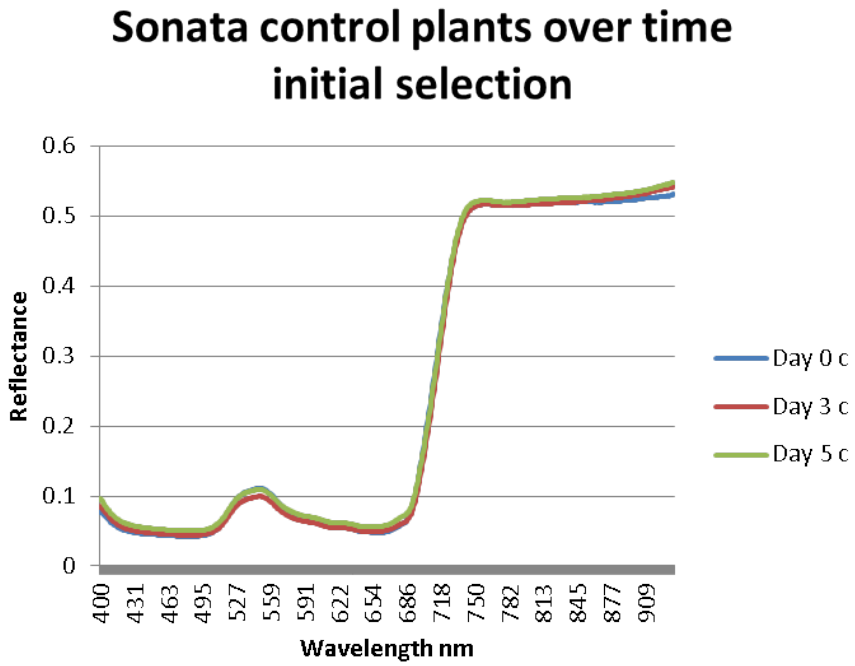

(a) The initial sonata plants spectral signatures over time.

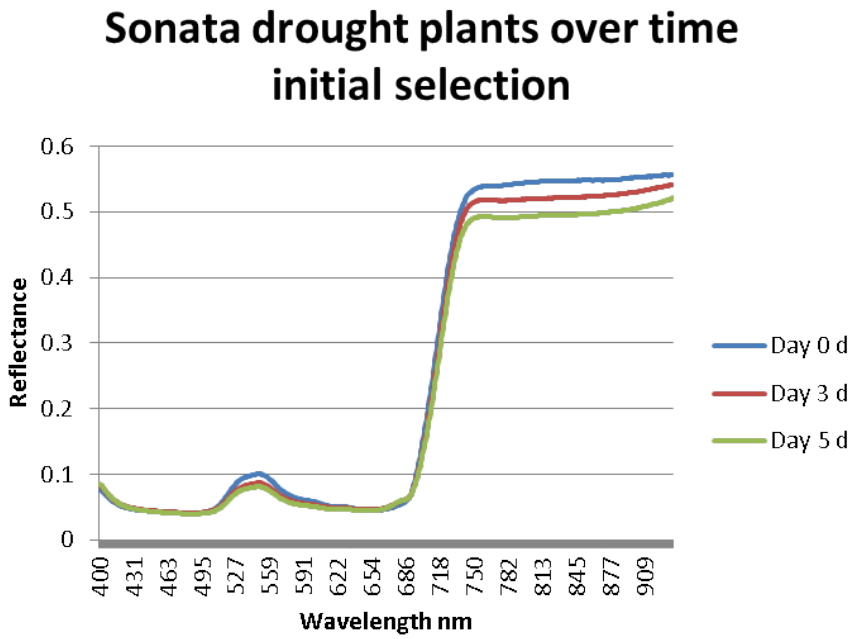

(b) The initial drought sonata plants spectral signatures over time.

Figure 7.2: (a) is the control plant graph from Chapter 3 presented for the initial analysis. (b) is the drought plant graph from Chapter 3 for the initial analysis.

The key questions are:

1. Can this manually measured result be replicated using the new automated

approach?

2. Does this improve the signal to noise ratio of the data (visible as more separation in the graphs)?

The average data from the 3 highest scoring leaves as calculated per earlier algorithm for all the Sonata plants have been plotted in Figure 7.3a and 7.3b. For Sonata, the control plants profiles appear similar over the time series. The drought plants spectral signatures have slightly decreased over the time series in the near infrared. The data has been smoothed using the Savitzky-Golay smoothing filter for 7points [117]:

$$y_t = (-2x_{t-3} + 3x_{t-2} + 6x_{t-1} + 7x_t + 6x_{t+1} + 3x_{t+2} - 2x_{t+3})/21 \tag{7.1}$$

During the early stage of water drought stress the hyperspectral signature decreases in the green and near-infrared sections of the spectrum. It is known that when the leaves dry out the red wavelengths will increase in reflectance because this is where the leaf pigments and chlorophyll affect the spectral signature. However during the early stages the green and near-infrared sections can decrease with little or no change to the red section. A study on Maize plants displays the spectral signatures of the leaves during different growth stages [118]. The spectral signatures are similar to the sonata drought spectral signatures, suggesting this is not an artifact.

The BSP14 control and drought plants' reflectivity increases slightly over time in Figure 7.4. The ordering effect is the same between drought and control, but ordering is different between Sonata drought and control (Figure 7.3). This might be explained, as BSP14 has some drought tolerance, as opposed to Sonata which does not. Alternatively this could be an anomaly with the BSP14 plants due to the change, since the control plants change more so than the Sonata plant data. The significance of these observations is hard to determine given the noise and variability in the datasets though.

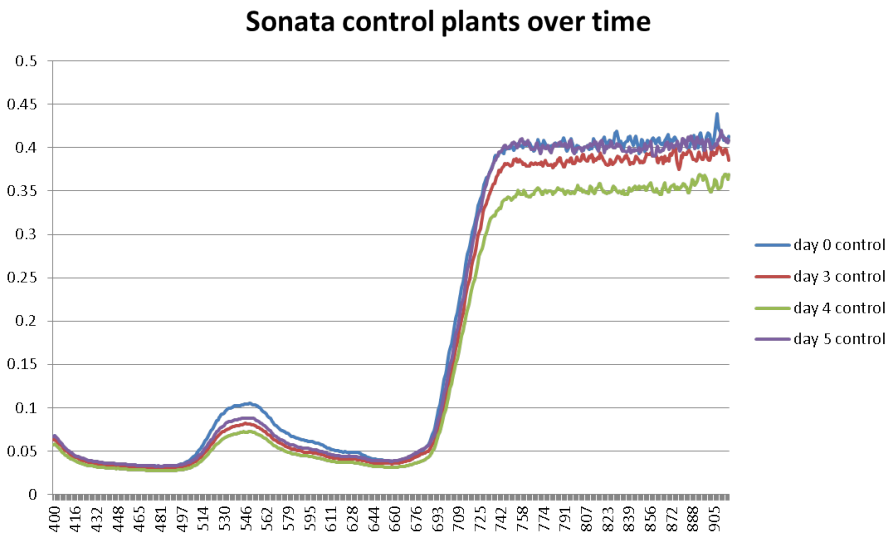

(a) The selected leaves for Sonata plants.

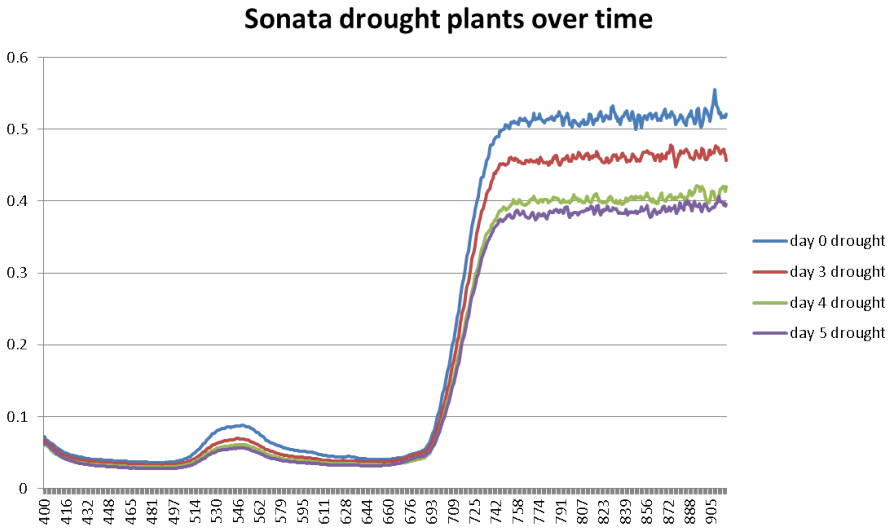

(b) The selected leaves for Sonata drought plants.

Figure 7.3: Leaves selected via the automated method introduced in Chapter 6. (a) The leaves for Sonata control plants over time. (b) The leaves for Sonata drought plants over time. The graphs have been smoothed using Savitzky-Golay.

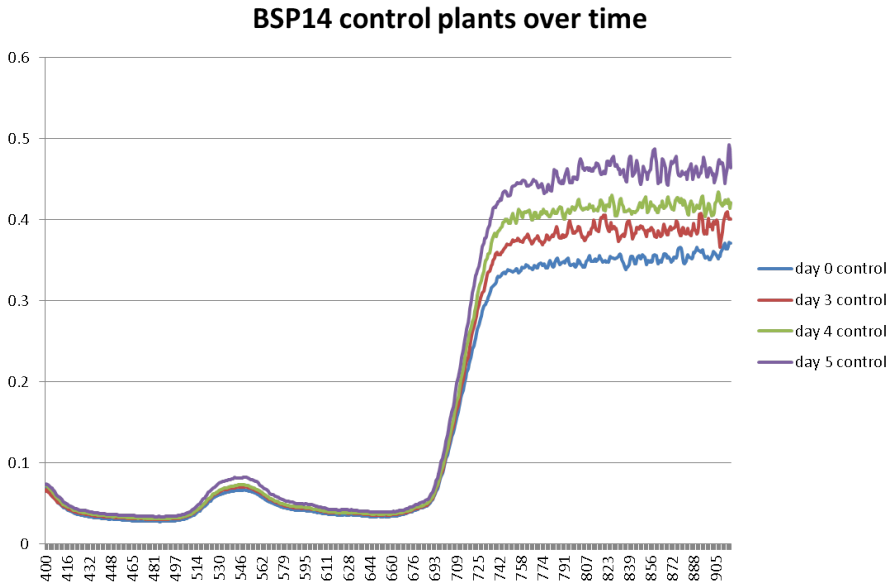

(a) The selected leaves for BSP14 plants.

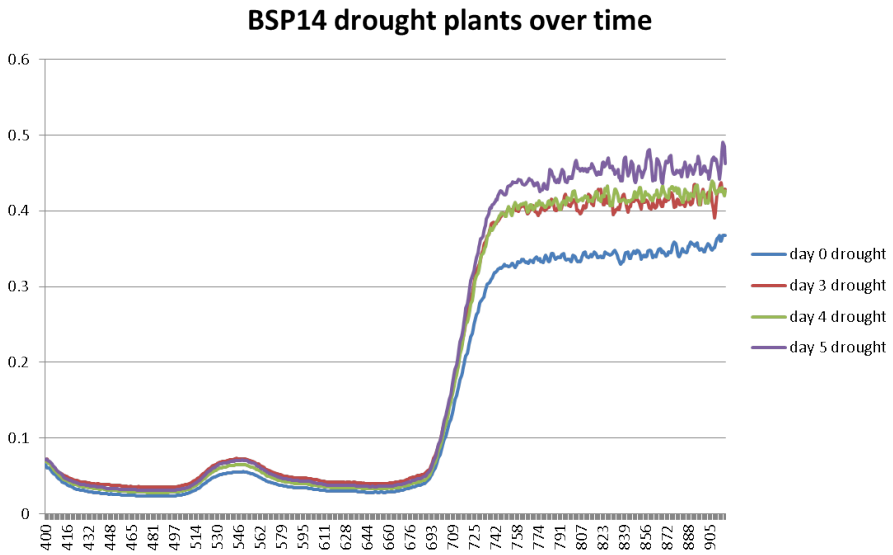

(b) The selected leaves for sonata drought plants.

Figure 7.4: Leaves selected via the automated method from Chapter 6 along with the same leaves for the time series. (a) The leaves for BSP14 control plants over time. (b) The leaves for BSP14 drought plants over time. The graphs have been smoothed using Savitzky-Golay.

7.3 Effect of normalisation

The effect of normalisation method is investigated to reveal further details from the dataset. The hyperspectral data presented so far is calibrated by dividing the

reflectance by the white balance [119]. This is how Figures 7.2,7.3 and 7.4 were created.

There are in fact many ways to further normalise the data and to investigate the effect of this [120]. A few methods have been used below. The first method is to divide the profile by a wavelength band using Equation 7.2 where R is the individual reflectance values, R660 is the reflectance at 660nm and NR is the normalised reflectance (See Figures 7.5 and 7.6). This normalises the data to a wavelength and the data will range from 1 to a maximum value. This helps to normalise profiles at one point in the spectral signatures but keep the possible changes visible in the graphs.

The wavelength is usually selected in the visible range where the data is less noisy. The wavelength selected is ideally a point which is not expected to change between datasets is used as a reference point. This normalises the brightness of reflectance at a particular colour.

$$NR = \frac{R}{R660} \tag{7.2}$$

The wavelength band selected varies if it is used in studies and there is no specific band (like NDVI).

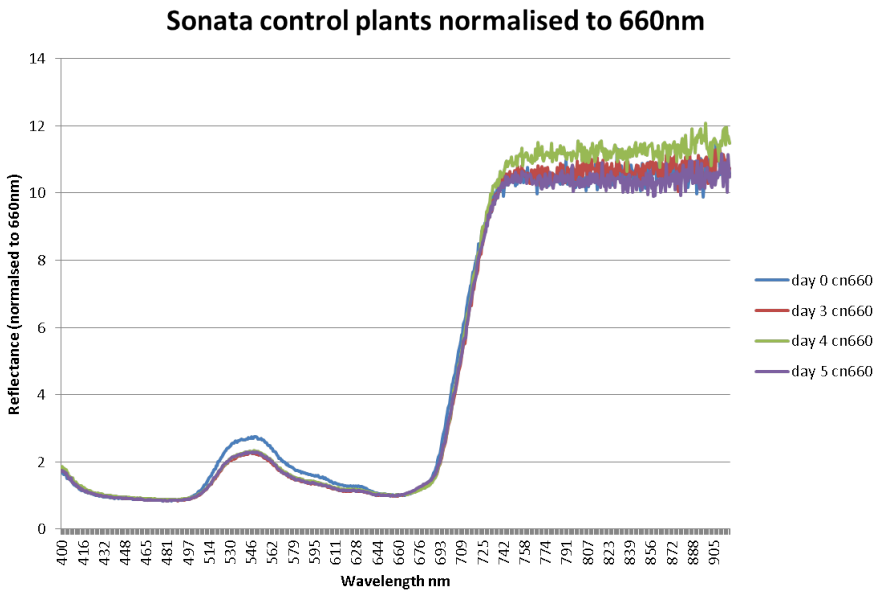

(a) Sonata control plants normalised at one wavelength - 660nm.

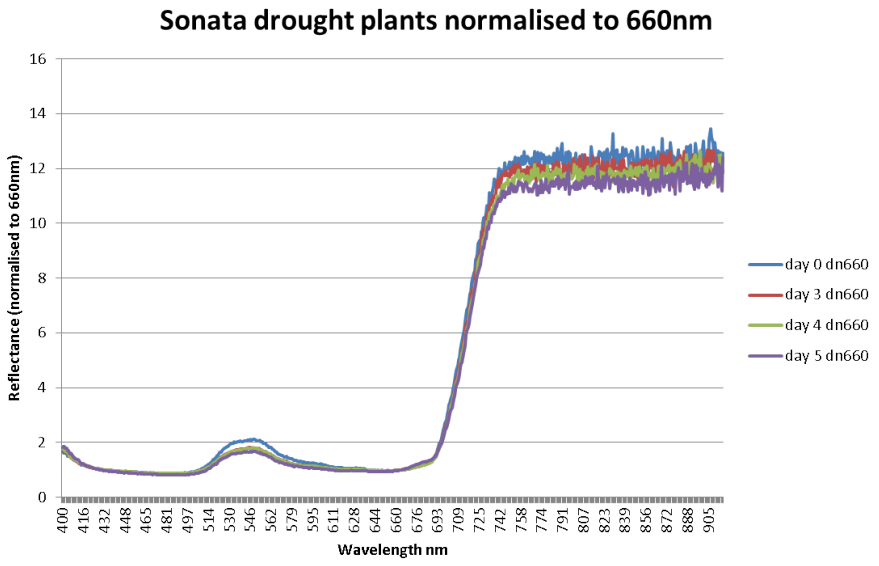

(b) Sonata drought plants normalised at 660nm.

Figure 7.5: The Sonata plants have been normalised to 1 at 660nm. This moves the wavelength to align at 660nm. A smaller but consistent seperation is still visible in the near infrared.

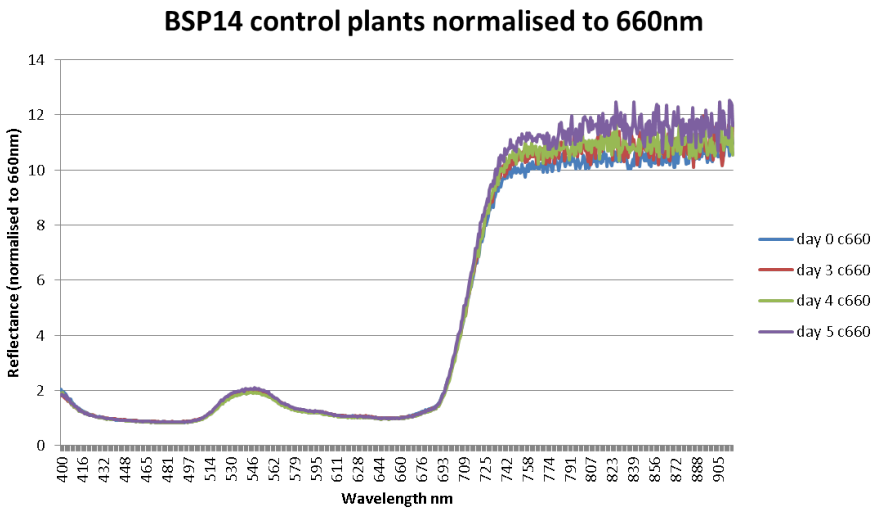

(a) BSP14 control plants normalised against one wavelength - 660nm.

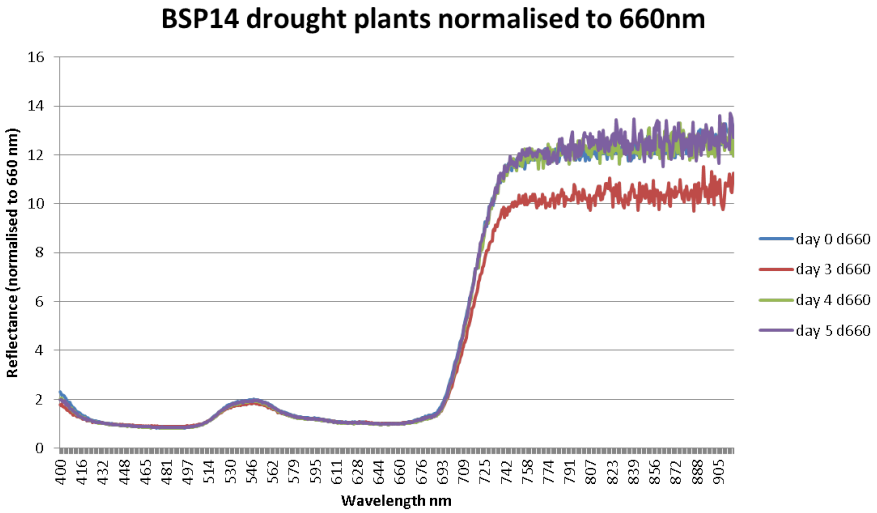

(b) BSP14 drought plants normalised against 660nm.

Figure 7.6: The BSP14 plants have been normalised to 1 at 660nm.

Figure 7.5b still shows a consistant decrease in the near infrared, but there is a smaller effect. Figures 7.5a and 7.6 appear to have no obvious effect.

A second normalisation process is to select the maximum and minimum reflectance values by Equation 7.3, where R is the individual reflectance, max is the maximum reflectance value, min is the minimum reflectance value and NR is the normalised reflectance. This normalises the data between 0 and 1, (see Figures 7.7 and 7.8).

$$NR = \frac{(R - min)}{(max - min)} \tag{7.3}$$

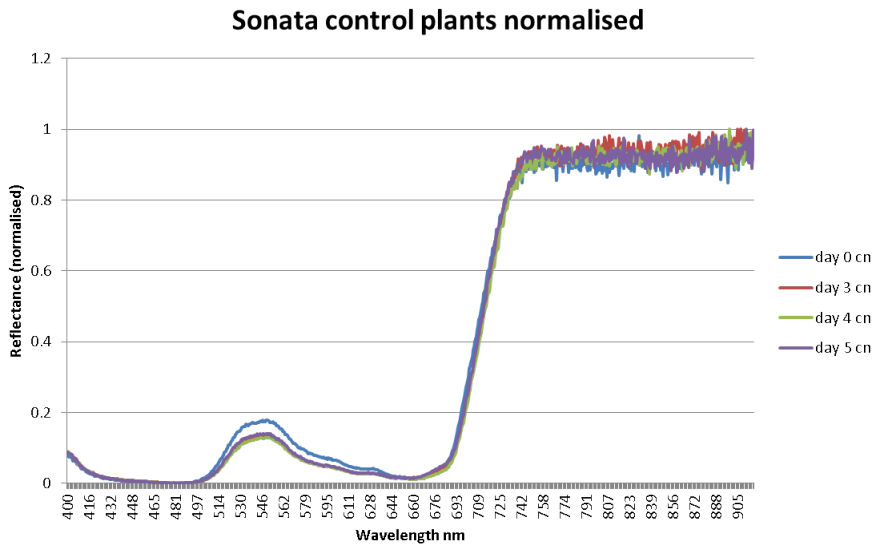

(a) Sonata control plants normalised in the 0-1 range.

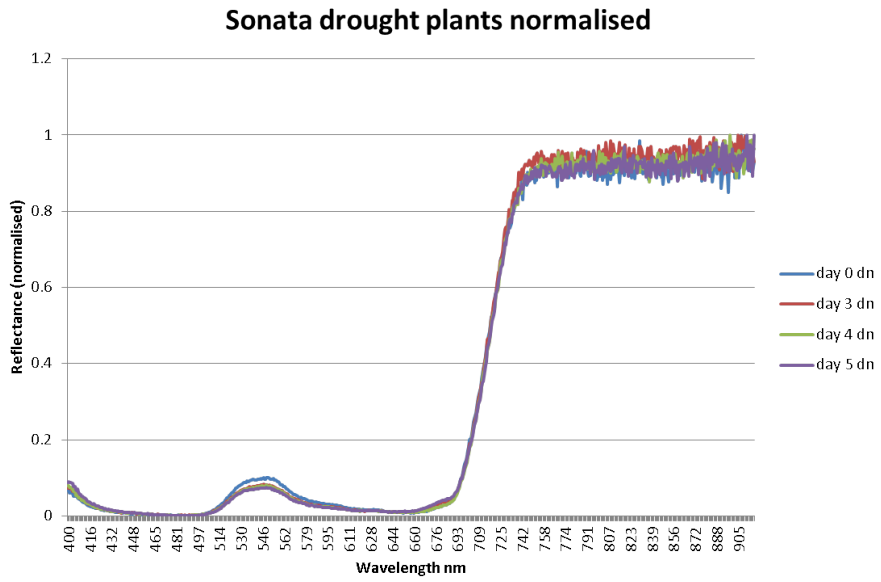

(b) Sonata drought plants normalised in the 0-1 range.

Figure 7.7: The Sonata plants have been normalised to the maximum and minimum wavelengths. While this normalises the data it is not the best option. The near-infrared section of the spectrum is quite noisy and the main change would appear in this section of the graph if it does.

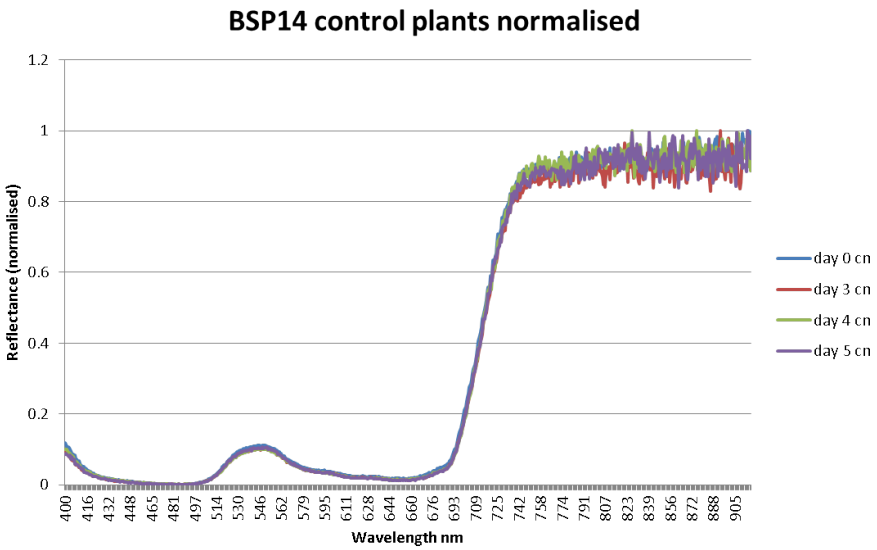

(a) BSP14 control plants normalised in the 0-1 range.

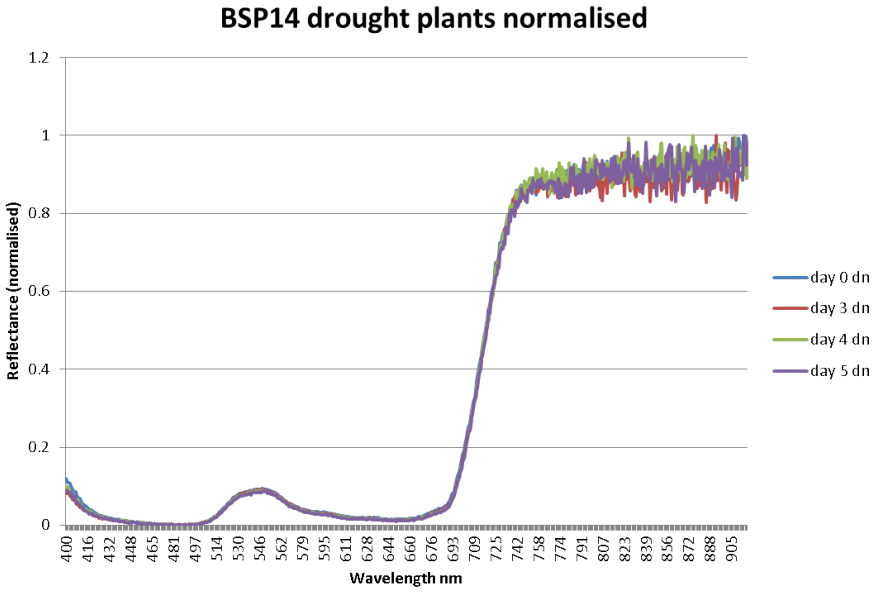

(b) BSP14 drought plants normalised in the 0-1 range.

Figure 7.8: The BSP14 plants have been normalised so the maximum and minimum wavelengths are 1 and 0 respectively.

The normalisation method using the maximum/minimum is sesitive to the minimum and maximum values which occur at noisy measurement points. However, no change detectable in BSP14, day 0 may be different in Sonata.

Conclusion of normalisation

Normalising might not be the best way to view hyperspectral data when the changes are being reduced or the data is being stretched where the change in data matters (for example the near infra-red and green wavebands). Normalisation helps to view the data in the same range. Several normalisation methods were applied to datasets for hyperspectral classification to see how well they would work. There are many ways to normalise data and the reflectance is usually 'normalised' against the white balance but there is not a single preferred method for normalising hyperspectral data yet.

Normalisation changes how data can be viewed. It might be sensitive to different effects, for example using the maximum/minimum method the data in the near infrared range appears the same however this is a key section for changes with drought stress. There needs to be a consideration before selecting a specific normalisation method for hyperspectral data.

7.4 Calculating indices from automatically selected leaves

Although complete profiles have been reported, here the effect of using this technique as an example vegetation index is examined.

Normalised Difference Vegetation Index (NDVI) is used to measure the health of crops, written as:

$$NDVI = \frac{(RNIR - RRED)}{(RNIR + RRED)} \tag{7.4}$$

This selects two wavelengths from the red (RRED) and near infrared (RNIR) ranges of the hyperspectral data and produces a value between 1 and -1 [56]. NDVI has been calculated from the data from Figures 7.3 and 7.4 to see if there is any change using the Vegetation index. The results shown are in Tables 7.1 and 7.2. For the Sonata Table there is a less than 0.01 decrease and the BSP14 Table varies over the time series. The NDVI results could not determine a difference during the short time period. The NDVI detects changes in the leaf colour which would change when wilting occurs. Most NDVI drought monitoring studies were used with remote sensing data and over long time periods; two separate studies

that used NDVI to compare it with different techniques noted that it took NDVI 13 and 16 days to indicate a difference in the plants [83] [84]. So not finding a positive result here is not surprising.

Table 7.1: NDVI for Sonata over the time series. This selects a wavelength in the near-infrared range and a wavelength in the red visible range

NDVI	Day 0	Day 3	Day 4	Day 5
Sonata drought	0.8545	0.8534	0.8459	0.8456
Sonata control	0.8254	0.8314	0.8323	0.8229

Table 7.2: NDVI for BSP14 over the time series. This selects a wavelength in the near-infrared range and a wavelength in the red visible range

NDVI	Day 0	Day 3	Day 4	Day 5
BSP14 drought	0.8450	0.8153	0.8464	0.8515
BSP14 control	0.8210	0.8229	0.8294	0.8407

NDVI can also be applied to the data per pixel and this has been applied to Sonata control and drought plants for the top 3 leaves selected.

Figure 7.9 displays the plants output from the NDVI. Figure 7.9a - 7.9b are the control and drought Sonata plants for day 3 and 7.9c and 7.9d are the control and drought plants for day 5. There does not appear to be much difference, NDVI is good for displaying changes over a longer period.

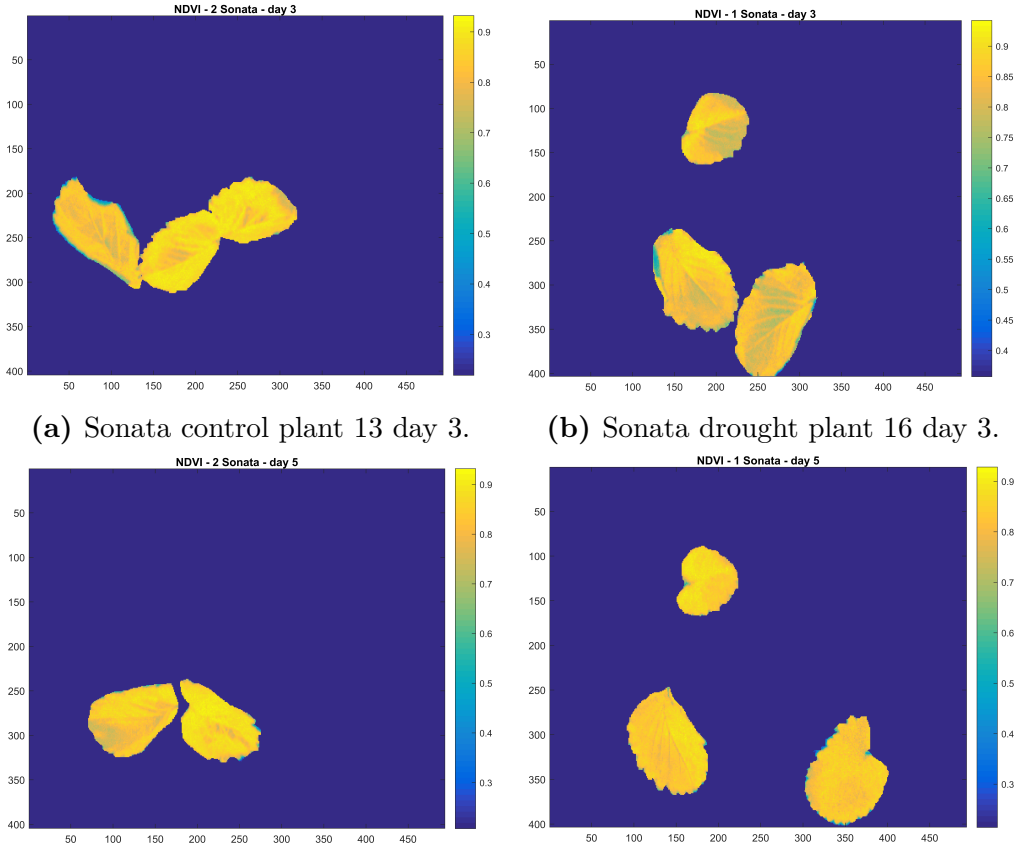

(a) Sonata control plant 13 day 3. **(b)** Sonata drought plant 16 day 3.

(c) Sonata control plant 13 day 5. **(d)** Sonata drought plant 16 day 5.

Figure 7.9: Two plants have been displayed over two time points. NDVI has been applied to the best 3 leaves. There does not appear to be much difference visible, NDVI is good for displaying changes over a longer period.

7.5 Conclusions

The graphs from this Chapter (Figures 7.3b and 7.5b) suggest that the leaf reflectance for Sonata decreases slightly in the near infrared range of the spectral signatures for the drought plants while the control plants remain almost the same, as indicated in the initial manual graphs. There is no difference between BSP control and drought plants observed (eg. Figure 7.4). Interesting, normalising removes some of the spread for BSP14 (Figure 7.6). This could indicate that the spread is from the brightness rather than colour. A few studies found that BSP is more resistant to drought than varieties NCC85-13V and Elvira. The paper also describes BSP14 as drought resistant. [88] [121]

While there appears to be a slight decrease in the infra red with the Sonata drought plants this is not enough to draw a conclusion that the drought can be detectable

on a specific day.

It is still challenging to identify a difference despite measuring from the same leaves over the time series (reducing the effects of leaf height and leaf angle changes). The best quality leaves were selected manually in Chapter 3 (2D) and automatically in Chapter 6 (3D), and still a clear conclusion cannot be drawn.

The difference between automatically selecting the best leaves over a time series and tracking the remaining time points results in a slight change in the leaves selected. The automatic method selected a slight variation over the time points due to leaves moving slightly and new leaves growing. This could mean when the leaves start to wilt the automatic method selects healthier/different leaves. However the tracked method selected the same leaves over the time series, this means consistency but also that the leaves could start to wilt.

Is picking different leaves good or bad? The two options are different approaches for selecting the leaves which would need to be considered depending on the effect under study.

The questions near the start of the Chapter can be answered:

1. Can the manually measured results from Chapter 3 be replicated using the new automated approach from Chapter 6? It appears that the manual observations have been automatically replicated.
2. Does this improve the signal to noise ratio of the data (visible as more of a separation in the graphs)? The data quality has not apparently been improved.

Chapter 8

Conclusions

This thesis developed a new approach to analyse hyperspectral data using a novel pipeline incorporating 3D data, with an aim to use hyperspectral data to determine if drought can be detected in strawberry plants before visible signs appear. This pipeline included 2D segmentation of leaves and 3D registration of hyperspectral data, and an approach to select suitable individual leaves from which to take hyperspectral measurements.

8.1 Summary

The problem was introduced in Chapter 1, and a review of relevant work presented in Chapter 2.

Chapter 3 explained the process of collecting the data, followed by calibration for the lighting and camera noise. Initial manual inspection was carried out on the data; leaves were selected manually, and data measured over 5 days to see if there was any difference between control and drought plants. The initial data collection and inspection indicated a difference in reflectance may be detectable over the range of days for drought but the data was noisy and changes were subtle so a new analysis approach was required.

Next, computer vision methods were developed to segment the individual leaves. This was developed in Chapter 4, using a leaf shape model. A shape template was translated to a leaf and rotated using the venetian pattern on the leaf before being incorporated as a shape prior into a level set. The results show that identified boundaries are of good quality, and of those segmented leaves, the Jaccard

accuracy was 0.68 for Sonata plants and 0.88 for BSP14 plants.

This method however is limited to just using the 2D spatial information. 3D information allows for the leaf shape, position and orientation to be taken into consideration when trying to automatically select the best leaves to use for analysis.

This leads to Chapter 5, which looked at using 3D information to support hyperspectral analysis. During the image capture process for the drought data, two extra sets of image data were captured. The first was an extra hyperspectral scan of each plant, but offset to the left of the first scan. This allowed for stereo data to be extracted. The second data collected were digital images from around a side view angle of the plant. Stereo did not produce good results, therefore multiview stereo was adopted.

First, the 3D models were built using existing multi view stereo software, PMVS, and surface reconstruction [112] [113]. Then the 3D model was projected to 2D, the hyperspectral data was registered with the 2D view then the hyperspectral data was mapped onto the 3D model using unique index values. The registration was compared in 2D, but evaluation was also carried out for the leaves in 3D. This 3D evaluation was to find out how well leaves were accurately mapped. The overall median values for the leaf mapping accuracy for Sonata was 83% and BSP14 was 75%, this only measured the percentage of correctly assigned leaf ID with the ground truth labelled leaf ID. There was hyperspectral data that did not exist for the model due to being cropped during imaging or misregistration, and hyperspectral data that could not be mapped onto the 3D model because small sections of the 3D model did not exist for the top down orientation of the plant.

Chapter 6 comprised of a novel high quality leaf selection process was developed. The best three leaves for each plant of the Sonata and BSP14 dataset were selected by finding the leaf angle, height and distance from the centre of the plant. Each element was scored from best to worst from the three measurements, and the best three leaves were selected from the combination of the three scores.

In Chapter 7 the same best three leaves were selected in the remaining time points to analyse the same leaf data over the time series. This was compared against the manual leaf selection plots from Chapter 3. While there appeared to be a slight decrease in the Sonata drought plants this was not enough to draw a conclusion that the drought could be detectable on a specific day.

It is still challenging to identify a difference despite measuring from the same leaves

over the time series (reducing the effects of leaf height and leaf angle changes). The best quality leaves were selected manually in Chapter 3 (2D) and automatically in Chapter 6 (3D), and still a clear conclusion cannot be drawn.

8.2 Challenges and Futurework

8.2.1 Full automation

In Chapter 5 the gross top-down pre-orientation of the 3D model was manually preset to save time (Eg solving a global registration problem) and then the proceeding local registration was automated. The camera setup was known and the geometry of the model was known.

Chapter 4 segmented leaves using a leaf model in 2D. Then Chapter 5 moved to using 3D models, and mapping the 2D hyperspectral data onto the model in order to find the orientation, location and general shape of the leaves to augment the corresponding hyperspectral data.

Currently the 2D ground truth is used to generate the leaf ID number; However it is possible for the level set segmentation from Chapter 4 to be inserted instead of the ground truth. In this case it would be fully automated. This has not been implemented in this case to focus on the ground truth and measuring how the data is mapped however is it possible.

After the best leaves were selected in Chapter 6 the same leaves were manually selected over the time series; this was manually selected to be sure that the leaves were the same through the time series. This could be automated by using a tracking algorithm. The plants were positioned in the same place and in the same orientation during imaging and the leaves should not move too far away from their original position (here, the day 3 images) until the stems weaken and wilting occurs therefore it should not be a challenging tracking problem for this time series data however if the leaves move it does increase the challenge of locating it. The level set leaf model (Chapter 4) produced the leaf boundaries and these could be used to track the same leaves over the time series, then compare the leaves on a pixel texture level to find the best match. However using the current leaf shape model a few leaves could potentially be missed because of the seed centre and vein location steps in the level set model.

8.2.2 Challenges of creating ground truth data

Two sets of ground truth data were created; the hyperspectral 2D ground truth and the 3D leaf ground truth.

The 2D ground truth consisted of manually drawing around the leaves in the image. The hyperspectral data was noisy but the leaf shapes were located. There were a few dark leaves near the bottom of plants where the boundary was estimated; this will always present a challenge for the manual creation of ground truth.

Manual creation of 2D ground truth is time consuming. A semi-automated approach could be used, such as level sets or other segmentation based methods, then any mis-segmented or completely missed leaves could be manually annotated afterwards.

The 3D ground truth consisted of manually assigning leaf IDs to the patches and triangles. This was challenging to create because the 3D model appeared as a collection of grey patches in Meshlab, there was no colour or texture information on the 3D model, only the patches and the shape they formed together to represent leaves. There were lots of patches and sometimes it was hard to see where the boundary of one leaf ended and the next leaf started. The SLR images were used as a visual reference to improve the leaf selection. BSP14 was in general easier to see where the leaves were. Sonata leaves were easy to annotate for the leaves on the top of the plant, however the lower leaves on the plant were hard to determine.

Summary

2D ground truth for this dataset could be semi-automated to an extent, although it would still require time to manually annotate leaves afterwards. 3D ground truth in this instance is very hard and would require manual annotation.

Ground truth is challenging and time consuming however it is required to evaluate how well the algorithm works. There are semi automated approaches that exist and this might become easier in the near future since deep learning can segment complex images on the same level of accuracy as a person can (ie. an expert).

8.2.3 Challenges of hyperspectral imaging

There are many challenges that arise with hyperspectral imaging.

When the plant is imaged in situ the leaves have different orientations to the camera plane, this can change the hyperspectral profile which can appear similar when looking for changes due to signs of stress. Therefore 3D information is needed to know the orientation of the leaves. With imaging a plant the leaves are at different heights and this also needs to be taken into consideration.

The hyperspectral data is calibrated with a white balance to account for the variance in lighting, however there is also a question of whether to normalise the data and if so, how. This was considered in Chapter 6.2; there are many methods to normalise data but are they helpful for hyperspectral data analysis and which one is the most appropriate? This must be considered in future work.

Hyperspectral data is a large, noisy dataset. The literature review in Chapter 2.2 discusses different approaches that can be used to analyse hyperspectral data during stress. A large portion of existing approaches reduce the data down to just several wavelengths to account for the size and noise of the data, but also to try and select important wavelengths.

There are a lot of factors to take into consideration with hyperspectral imaging and challenges that need to be overcome.

8.2.4 Application to other stresses

The work in this thesis used the drought dataset, a logical question is if this could be extended to other stresses. The methods in Chapters 4 and 5 can be expected to work if the datasets are similar (eg. hyperspectral images and high resolution multi-view images). Chapter 6.1 selects the leaf measurements, therefore it could be applied to a mildew dataset because the leaf angles and the leaf height would still be important.

This thesis has created a hyperspectral dataset for three plant stresses over a time series for early onset of stress until visible signs appeared (Objective 1). The leaf measurements were manually selected and analysed to see a small change in the near infrared range of the data for the drought times series while the control times series remained consistent (Objective 2). A leaf shape model was developed to automatically locate and segment individual leaves (in 2D) with good results (Objective 3). Then a pipeline was developed to use 3D information to improve the hyperspectral measurements (Objective 4). The hyperspectral measurements were used to find the best leaves on the plant using the location and inclination

of the leaves (Objective 5). The same automated leaves were manually selected for the other days to compare with the initial manual analysis from Chapter 3 (Objective 6).

A complete, largely automated pipeline has been developed, which incorporated hyperspectral and 3D data.

Final results suggest a possible difference in reflectance for some varieties in early stages of drought but a firm conclusion is hard to come to in part because of the challenges with the variables present in hyperspectral imaging. There are studies that have managed to detect early stress in comparison to vegetation indices, in certain crops, using hyperspectral imaging however the data can be noisy and the analysis approach needs to be considered.

Appendix A

Seed evaluation for chapter 4

Table A.1: BSP14 plants 1 and 12 individual leaf ground truth and seed centre points with the x and y differences. The units are measured in pixels.

Plant	Leaf	Gt x	Gt y	Seed x	Seed y	Distance	X (+-)	Y (+-)
1	1	603.475	114.5639	608.6849	114.3853	5.212975	5.209912	-0.17866
1	3	416.5448	106.7693	410.3397	103.2367	7.140121	-6.20505	-3.53252
1	5	430.5354	157.8049	435.3597	159.7593	5.205096	4.82425	1.954389
1	9	490.0699	235.1528	473.6718	228.7836	17.5916	-16.3981	-6.36923
1	10	430.1237	300.6735	432.3587	300.379	2.254339	2.235021	-0.29449
1	11	520.6388	311.2736	523.443	309.9622	3.095753	2.804256	-1.31142
1	13	539.954	146.8688	536.0079	145.7048	4.114163	-3.94608	-1.16396
1	20	537.8276	51.62356	534.947	49.53851	3.556027	-2.88061	-2.08505
12	1	532.5509	317.2531	529.9828	320.9776	4.524114	-2.56811	3.724571
12	4	613.434	196.7593	616.0307	192.7978	4.736679	2.596729	-3.96146
12	5	547.9732	180.9703	544.4433	180.003	3.659968	-3.52983	-0.96731
12	9	437.2439	164.0809	437.7388	162.1506	1.992795	0.494881	-1.93037
12	10	410.948	248.2226	415.5051	250.6378	5.157589	4.557137	2.415209
12	11	340.1932	182.1726	339.982	180.8359	1.353318	-0.21123	-1.33673
12	12	426.4942	93.66181	426.2491	85.94624	7.719468	-0.24507	-7.71558
12	14	533.0788	38.39824	534.5282	48.1066	9.815964	1.449444	9.70836
12	16	613.5257	90.83576	614.1727	92.69697	1.970476	0.647048	1.86121
12	19	618.8872	335.2881	620.9139	340.1288	5.247824	2.026686	4.840682

Table A.2: Sonata plants 13 and 14 individual leaf ground truth and seed centre points with the x and y differences. This is after duplicate seeds have been removed. The units are measured in pixels.

Plant	Leaf	Gt x	Gt y	Seed x	Seed y	Distance	X (+-)	Y (+-)
13	1	363.2285	307.2511	359.7379	314.5678	8.10669	-3.49067	7.316669
13	2	397.285	374.1244	393.4149	398.0421	24.22878	-3.8701	23.9177
13	3	467.304	338.7141	488.4514	378.2597	44.84493	21.14739	39.54561
13	4	525.6484	222.0119	527.1525	216.2567	5.948441	1.504007	-5.75516
13	5	563.8384	297.3112	548.0503	308.0217	19.0782	-15.7881	10.71043
13	6	638.2471	264.632	648.5584	281.2319	19.54174	10.31131	16.59989
13	7	538.794	123.8247	534.62	126.4413	4.926321	-4.174	2.616553
13	8	611.819	98.3948	620.8402	111.4963	15.90696	9.021151	13.10153
13	11	431.428	212.0348	417.0442	215.7668	14.86007	-14.3838	3.732004
13	13	674.2087	184.2801	682.0637	182.1762	8.131809	7.854949	-2.10383
13	20	472.3487	28.04273	466.056	26.38404	6.50758	-6.29264	-1.65869
13	21	587.3517	31.244	613.511	27.06611	26.49081	26.15928	-4.17789
14	3	592.6471	249.6762	597.4975	262.886	14.07217	4.850402	13.20983
14	4	490.9494	203.911	457.4945	215.3479	35.35574	-33.4548	11.43689
14	5	582.371	137.6175	585.0355	133.4946	4.909046	2.664581	-4.12295
14	7	435.4485	304.1413	431.2114	317.818	14.31803	-4.23709	13.67674
14	11	519.2559	71.46565	528.855	58.64546	16.0156	9.599079	-12.8202
14	15	650.9598	70.06473	677.8204	56.24876	30.20549	26.86058	-13.816

Appendix B

Protocols - from Chapter 3

B.1 Strawberry Plants - Mildew experiment protocol

Plant Material Preparation

Dormant strawberry plants were purchased from strawberry propagators and kept in a cold store at 2°C until required. The dormant strawberry plants were taken out of the cold store the night before potting and left at +2°C to ensure plants were not damaged due to a sudden change in temperature. The following day, the strawberry plants were potted into FP9 pots using standard compost and grown for 2-3 weeks until they have 2-3 leaves. When plants have 2-3 leaves, a thorough inspection of the leaves was undertaken and any marks/signs on the leaves of any other diseases/pests were noted. A dissecting microscope is used to check the leaves on 10% of the plants at random to ensure they were pest-free.

Equipment: FP9 pots, Standard compost, Cold store plants, Glasshouse for the plants to grow in.

Location for the plants during the experiment The experimental strawberry plants were grown in CE cabinets at NIAB EMR. The CE cabinets were cleaned thoroughly with 70% ethanol the day before the experiment to make sure no disease/pests were present that would infect the control/infected plants during the experiment. After the experiment was finished, the cabinets were cleaned again using 70% ethanol.

Mildew Infection process

Two upside-down plant saucers were placed into the CE cabinet ready for the large mildew infected plants to be placed upon. Previously infected mildew plants were transferred from the glasshouse to the CE cabinets. The plants were transferred in a completely covered box to prevent spreading of mildew spores to any nearby strawberry plants. The two large mildew plants were placed into the middle of cabinet but evenly spaced. The test strawberry plants, cv. 'Elsanta'and 'Fenella'were placed in a circle around the two plants to evenly expose them to the mildew spores.

Equipment: Mildew infected strawberry plants (x2), Box with lid and box, 25 plants (11 Elsanta plants and 14 Fenella plants)

7 Fenella plants and 5 Elsanta plants will be placed in Cabinet 6 as control plants. 7 Fenella plants and 6 Elsanta plants will be placed in Cabinet 7 with the two mildew plants.

Imaging Process

The strawberry plants were imaged before they were placed into the cabinets (day 0). They were then imaged every other day until there were visible signs on all infected plants. To transport the plants from the cabinets for imaging, a box with a lid was used to contain the mildew spores. Water the plants every other day. Also after each imaging session the plants are placed back into random positions.

Imaging method at EMR

1. Clean the area around the camera with 70% ethanol
2. Turn on the three switches at the plugs and turn the PC on.
3. Open the HSI software and select 'start pipeline'
4. Switch the lights on at the voltage box, bottom left and then bottom right.
5. Move the board using the arrow keys on the screen if it is needed.
6. Turn the main room light off.
7. Place the white balance on the board.
8. Select the white balance calibration button
9. Cover lens with hand and select dark (white) calibration button
10. Cover lens with hand and select dark (object) calibration button
11. Place the plant on the board under the lens.
12. Make sure the lens is in focus.
13. Move the board to the desired start position using the arrow keys on the screen

(Repeat steps 14-20 for each plant)

14. Take a digital image of the plant
15. Place the plant on the board
16. Click the right arrow key on screen to take the board to the start position
17. (select right arrow key to let the plants speed be tested, the image on the screen needs to match on the left and right sides)
18. Change the number of the speed slightly and repeat the last three steps until the image matches on the screen.
19. Name the file a relevant name ie. Day0_Fenella1_Cal
20. Select frame arrow right key on screen to finally capture the images.

B.2 Strawberry Plants - Spider mite experiment protocol

Plant Material Preparation

Dormant strawberry plants were purchased from strawberry propagators and kept in a cold store at 2ÂřC until required. The dormant strawberry plants were taken out of the cold store the night before potting and left at +2 °C to ensure plants were not damaged due to a sudden change in temperature. The following day, the strawberry plants were potted into FP9 pots using standard compost and grown for 2-3 weeks until they have 2-3 leaves. When plants have 2-3 leaves, a thorough inspection of the leaves was undertaken and any marks/signs on the leaves of any other diseases/pests were noted. A dissecting microscope is used to check the leaves on 10% of the plants at random to ensure they were pest-free.

Equipment:

- FP9 pots
- Standard compost
- Cold store plants
- Glasshouse for the plants to grow in
- Number the plants for the experiment 1-10 for the control and 11-20 for the inoculated.

Location for the plants during the experiment

The plants were kept in a contained glasshouse that limits exposure of spider mites. The area was cleaned before and after the experiment using 70% ethanol.

Spider mite preparation

Use spider mite infected plant. Remove and cut the leaf into sections using a microscope and put the sections into a sealed container. Do this on the day that the plants will be infected. Set up the area in the glasshouse with a tray filled with water and place 10 saucers upside down onto the tray then put the plants on these saucers. This creates a barrier and keeps the spider mites location as controlled as possible.

Infection process

Place an infected leaf section on the bottom of the stem of each plant. This will allow an even distribution and also a natural infection compared to placing the infected leaf on the top of a leaf.

Equipment:

- Spider mite leaf sections (x10)
- 10 plants (5 per variety) that will be infected.
- 2 (or 3 if possible) varieties.

Imaging Process Image before the plants are placed into the glasshouse (day 0). Then image every day (10am-2pm) until there are visible signs on the plants.

Imaging method at EMR

1. Clean the area around the camera with 70% ethanol
2. Turn on the three switches at the plugs and turn the PC on.
3. Open the HSI software and select 'start pipeline'
4. Switch the lights on at the voltage box, bottom left and then bottom right.
5. Move the board using the arrow keys on the screen if it is needed.
6. Turn the main room light off.
7. Place the white balance on the board.
8. Select the white balance calibration button
9. Cover lens with hand and select dark (white) calibration button
10. Cover lens with hand and select dark (object) calibration button
11. Place the plant on the board under the lens.
12. Make sure the lens is in focus.

13. Move the board to the desired start position using the arrow keys on the screen

    (Repeat steps 14-20 for each plant)
14. Take a digital image of the plant
15. Place the plant on the board
16. Click the right arrow key on screen to take the board to the start position
17. (select right arrow key to let the plants speed be tested, the image on the screen needs to match on the left and right sides)
18. Change the number of the speed slightly and repeat the last three steps until the image matches on the screen.
19. Name the file a relevant name ie. Day0_Fenella1_Cal
20. Select frame arrow right key on screen to finally capture the images.

Recording additional information

- plant variety, number, date, time and time after infection
- Speed, start and stop position, lighting, exposure rate and frame rate
- Any visible marks/signs (including before infection to discount it)
- General health
- Grading 1-4 where 1 is no signs, 2 is slight marks or discoloration, 3 is a few light green speckles and 4 is obvious light green speckles

Appendix C

Individual leaves with the mapped data - for Chapter 5

Individual leaves have been selected to display the hyperspectral data mapped onto the leaf in Figure C.1.

(a) Plant 12 leaf 2 gt

(b) Plant 12 leaf 2 map.

(c) Plant 12 leaf 3 gt

(d) Plant 12 leaf 3 map.

(e) Plant 12 leaf 12 gt

(f) Plant 12 leaf 12 map.

Figure C.1: Plant 12 leaf comparison.

Bibliography

[1] Horticulture statistics 2015, . https://www.gov.uk/government/statistics/horticulture-statistics-2015.

[2] Horticulture statistics 2016, . https://www.gov.uk/government/statistics/horticulture-statistics-2016.

[3] C Swift. Strawberry diseases. http://www.ext.colostate.edu/pubs/garden/02931.html, 2011. Accessed: 2015-02-16.

[4] Emanuele Livinali, Raul Antonio Sperotto, Noeli Juarez Ferla, and Claucia Fernanda Volken de Souza. Physicochemical and nutritional alterations induced by two-spotted spider mite infestation on strawberry plants. *Electronic Journal of Biotechnology*, 17(5):193–198, 2014.

[5] Stijn Dhondt, Nathalie Wuyts, and Dirk Inzé. Cell to whole-plant phenotyping: the best is yet to come. *Trends in plant science*, 18(8):428–439, 2013.

[6] Sotirios A Tsaftaris and Christos Noutsos. Plant phenotyping with low cost digital cameras and image analytics. In *Information Technologies in Environmental Engineering*, pages 238–251. Springer, 2009.

[7] Gerie van der Heijden, Yu Song, Graham Horgan, Gerrit Polder, Anja Dieleman, Marco Bink, Alain Palloix, Fred van Eeuwijk, and Chris Glasbey. Spicy: towards automated phenotyping of large pepper plants in the greenhouse. *Functional Plant Biology*, 39(11):870–877, 2012.

[8] H Genc, L Genc, H Turhan, SE Smith, and JL Nation. Vegetation indices as indicators of damage by the sunn pest (hemiptera: Scutelleridae) to field grown wheat. *African Journal of Biotechnology*, 7(2), 2008.

[9] J Verrelst, B Koetz, M Kneubühler, and M Schaepman. Directional sensitivity analysis of vegetation indices from multi-angular chris/proba data. In *ISPRS Commission VII Mid-term symposium*, pages 677–683, 2006.

[10] Ribana Roschera, Jan Behmanna, Anne-Katrin Mahleinc, Jan Dupuisa, Heiner Kuhlmanna, and Lutz Plümera. Detection of disease symptoms on hyperspectral 3d plant models. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 89–96, 2016.

[11] Alexandre Y Fong and Elliot Wachman. Hyperspectral imaging for the life sciences. *Biophotonics International*, 15(3):38, 2008.

[12] Oi Wah Liew, Pek Ching Jenny Chong, Bingqing Li, and Anand K Asundi. Signature optical cues: emerging technologies for monitoring plant health. *Sensors*, 8(5):3205–3239, 2008.

[13] Amy Lowe, Nicola Harrison, and Andrew P French. Hyperspectral image analysis techniques for the detection and classification of the early onset of plant disease and stress. *Plant Methods*, 13(1):80, 2017.

[14] The grower oct/nov 2017 ahdb horticulture. `https://horticulture.ahdb.org.uk/publication/grower-octnov-2017`.

[15] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6):679–698, 1986.

[16] Irwin Sobel and Gary Feldman. A 3x3 isotropic gradient operator for image processing. *a talk at the Stanford Artificial Project in*, pages 271–272, 1968.

[17] Suman Tatiraju and Avi Mehta. Image segmentation using k-means clustering, em and normalized cuts. *University Of California Irvine*, 2008.

[18] Keinosuke Fukunaga and Larry Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *Information Theory, IEEE Transactions on*, 21(1):32–40, 1975.

[19] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(5):603–619, 2002.

[20] Xin Huang and Liangpei Zhang. An adaptive mean-shift analysis approach for object extraction and classification from urban hyperspectral imagery.

Geoscience and Remote Sensing, IEEE Transactions on, 46(12):4173–4185, 2008.

[21] Dorothy M Greig, Bruce T Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989.

[22] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8): 888–905, 2000.

[23] N Gorretta, JM Roger, G Rabatel, V Bellon-Maurel, C Fiorio, and C Lelong. Hyperspectral image segmentation: the butterfly approach. In *Hyperspectral Image and Signal Processing: Evolution in Remote Sensing, 2009. WHISPERS'09. First Workshop on*, pages 1–4. IEEE, 2009.

[24] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988.

[25] Stanley Osher and James A Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of computational physics*, 79(1):12–49, 1988.

[26] Ravi Malladi, James A Sethian, and Baba C Vemuri. Shape modeling with front propagation: A level set approach. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 17(2):158–175, 1995.

[27] Tony F Chan and Luminita A Vese. Active contours without edges. *Image processing, IEEE transactions on*, 10(2):266–277, 2001.

[28] Nirmal Keshava and John F Mustard. Spectral unmixing. *Signal Processing Magazine, IEEE*, 19(1):44–57, 2002.

[29] Nirmal Keshava. A survey of spectral unmixing algorithms. *Lincoln Laboratory Journal*, 14(1):55–78, 2003.

[30] Wei Tang, Zhenwei Shi, and Zhenyu An. Nonnegative matrix factorization for hyperspectral unmixing using prior knowledge of spectral signatures. *Optical Engineering*, 51(8):087001–1, 2012.

[31] José MP Nascimento and José M Bioucas-Dias. Dependent component analysis: A hyperspectral unmixing algorithm. In *Pattern Recognition and Image Analysis*, pages 612–619. Springer, 2007.

[32] Roozbeh Rajabi and Hassan Ghassemian. Spectral unmixing of hyperspectral imagery using multilayer nmf. 2014.

[33] Chein-I Chang and Qian Du. Estimation of number of spectrally distinct signal sources in hyperspectral imagery. *Geoscience and Remote Sensing, IEEE Transactions on*, 42(3):608–619, 2004.

[34] Stefan M Soltuz, Wenwu Wang, and Philip JB Jackson. A hybrid iterative algorithm for nonnegative matrix factorization. In *Statistical Signal Processing, 2009. SSP'09. IEEE/SP 15th Workshop on*, pages 409–412. IEEE, 2009.

[35] O Berné, A Boulais, Y Deville, C Joblin, J Pety, D Teyssier, and M Gerin. Nmf-based spectral unmixing of herschel observations of the horsehead nebula. *Proceedings of the 5th Workshop on Hyperspectral Image and Signal Processing (WHISPERS 2013), Gainesville, Florida, USA*, 2013.

[36] Sen Jia and Yuntao Qian. Constrained nonnegative matrix factorization for hyperspectral unmixing. *Geoscience and Remote Sensing, IEEE Transactions on*, 47(1):161–173, 2009.

[37] Hanno Scharr, Massimo Minervini, Andrew P. French, Christian Klukas, David M. Kramer, Xiaoming Liu, Imanol Luengo, Jean-Michel Pape, Gerrit Polder, Danijela Vukadinovic, Xi Yin, and Sotirios A. Tsaftaris. Leaf segmentation in plant phenotyping: a collation study. *Machine Vision and Applications*, 27(4):585–606, 2016.

[38] Danijela Vukadinovic and Gerrit Polder. Watershed and supervised classification based fully automated method for separate leaf segmentation. In *The Netherland Congress on Computer Vision*, pages 1–2, 2015.

[39] Kyle Simek and Kobus Barnard. Gaussian process shape models for bayesian segmentation of plant leaves. *Proceedings of the Computer Vision Problems in Plant Phenotyping (CVPPP), pages*, pages 4–1, 2015.

[40] Long Quan, Ping Tan, Gang Zeng, Lu Yuan, Jingdong Wang, and Sing Bing Kang. Image-based plant modeling. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 599–604. ACM, 2006.

[41] Xiao-Feng Wang and Hai Min. An efficient two-stage level set segmentation framework for overlapping plant leaf image. In *Intelligent Computing Technology*, pages 466–474. Springer, 2012.

[42] Mengye Ren and Richard S Zemel. End-to-end instance segmentation with recurrent attention. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA*, pages 21–26, 2017.

[43] Bernardino Romera-Paredes and Philip Hilaire Sean Torr. Recurrent instance segmentation. In *European Conference on Computer Vision*, pages 312–329. Springer, 2016.

[44] Bert De Brabandere, Davy Neven, and Luc Van Gool. Semantic instance segmentation with a discriminative loss function. In *Deep Learning for Robotic Vision, workshop at CVPR 2017*, pages 1–2. CVPR, 2017.

[45] Andrei Dobrescu, Mario Valerio Giuffrida, and Sotirios A Tsaftaris. Leveraging multiple datasets for deep leaf counting. In *Computer Vision Workshop (ICCVW), 2017 IEEE International Conference on*, pages 2072–2079. IEEE, 2017.

[46] Serge Savary, Andrea Ficke, Jean-Noel Aubertot, and Clayton Hollier. Crop losses due to diseases and their implications for global food production losses and food security. *Food Security*, 4(4):519–537, July 2012. ISSN 1876-4517, 1876-4525. doi: 10.1007/s12571-012-0200-5. URL `http://link.springer.com/article/10.1007/s12571-012-0200-5`.

[47] E.-C. Oerke. Crop losses to pests. *The Journal of Agricultural Science*, 144(01):31–43, February 2006. ISSN 1469-5146. doi: 10.1017/S0021859605005708.

[48] M. B. Riley, M. R. Williamson, and O. Maloy. Plant disease diagnosis. *The Plant Health Instructor*, 2002. ISSN 1935-9411. doi: 10.1094/PHI-I-2002-1021-01. URL `http://www.apsnet.org/edcenter/intropp/topics/Pages/PlantDiseaseDiagnosis.aspx`.

[49] Anne-Katrin Mahlein. Plant disease detection by imaging sensors–parallels and specific demands for precision agriculture and plant phenotyping. *Plant Disease*, 100(2):241–251, 2016.

[50] Josep Penuelas and Iolanda Filella. Visible and near-infrared reflectance techniques for diagnosing plant physiological status. *Trends in plant science*, 3(4):151–156, 1998. URL http://www.sciencedirect.com/science/article/pii/S1360138598012138.

[51] Melvin B. Satterwhite and J. P. Henley. Hyperspectral Signatures (400 to 2500 nm) of Vegetation, Minerals, Soils, Rocks, and Cultural Features: Laboratory and Field Measurements. Technical report, December 1990.

[52] J. W. Rouse Jr. Monitoring the vernal advancement and retrogradation (green wave effect) of natural vegetation. 1972. URL http://ntrs.nasa.gov/search.jsp?R=19730009607.

[53] Rosa Lasaponara and Nicola Masini. Detection of archaeological crop marks by using satellite QuickBird multispectral imagery. *Journal of Archaeological Science*, 34(2):214–221, February 2007. ISSN 0305-4403. doi: 10.1016/j.jas.2006.04.014. URL http://www.sciencedirect.com/science/article/pii/S0305440306000951.

[54] A. K. Mahlein, T. Rumpf, P. Welke, H. W. Dehne, L. PlÃijmer, U. Steiner, and E. C. Oerke. Development of spectral indices for detecting and identifying plant diseases. *Remote Sensing of Environment*, 128:21–30, January 2013. ISSN 0034-4257. doi: 10.1016/j.rse.2012.09.019. URL http://www.sciencedirect.com/science/article/pii/S0034425712003793.

[55] Davoud Ashourloo, Mohammad Reza Mobasheri, and Alfredo Huete. Developing Two Spectral Disease Indices for Detection of Wheat Leaf Rust (Pucciniatriticina). *Remote Sensing*, 6(6):4723–4740, May 2014. doi: 10.3390/rs6064723.

[56] ENVI. Vegetation analysis: Using vegetation indices in envi. URL http://www.exelisvis.com/Learn/WhitepapersDetail/TabId/802/ArtMID/2627/ArticleID/13742/Vegetation-Analysis-Using-Vegetation-Indices-in-ENVI.aspx.

[57] Daniel A Sims and John A Gamon. Relationships between leaf pigment content and spectral reflectance across a wide

range of species, leaf structures and developmental stages. *Remote Sensing of Environment*, 81(2âĂŞ3):337–354, August 2002. ISSN 0034-4257. doi: 10.1016/S0034-4257(02)00010-X. URL http://www.sciencedirect.com/science/article/pii/S003442570200010X.

[58] J. Penuelas, I. Filella, P. Lloret, F. Munoz, and M. Vilajeliu. Reflectance assessment of mite effects on apple trees. *International Journal of Remote Sensing*, 16(14):2727–2733, 1995. URL http://www.tandfonline.com/doi/abs/10.1080/01431169508954588.

[59] J. Penuelas, F. Baret, and I. Filella. Semi-empirical indices to assess carotenoids/chlorophyll a ratio from leaf spectral reflectance. *Photosynthetica*, 31:221–230, 1995.

[60] Moses Azong Cho and Andrew K. Skidmore. A new technique for extracting the red edge position from hyperspectral data: The linear extrapolation method. *Remote Sensing of Environment*, 101(2):181–193, March 2006. ISSN 0034-4257. doi: 10.1016/j.rse.2005.12.011.

[61] Xueren Cao, Yong Luo, Yilin Zhou, Jieru Fan, Xiangming Xu, Jonathan S. West, Xiayu Duan, and Dengfa Cheng. Detection of Powdery Mildew in Two Winter Wheat Plant Densities and Prediction of Grain Yield Using Canopy Hyperspectral Reflectance. *PLOS ONE*, 10(3):e0121462, March 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0121462.

[62] Cedric Bravo, Dimitrios Moshou, Jonathan West, Alastair McCartney, and Herman Ramon. Early Disease Detection in Wheat Fields using Spectral Reflectance. *Biosystems Engineering*, 84(2):137–145, February 2003. ISSN 1537-5110. doi: 10.1016/S1537-5110(02)00269-6.

[63] Dimitrios Moshou, Cedric Bravo, Jonathan West, Stijn Wahlen, Alastair McCartney, and Herman Ramon. Automatic detection of yellow rust in wheat using reflectance measurements and neural networks. *Computers and Electronics in Agriculture*, 44(3):173–188, September 2004. ISSN 0168-1699. doi: 10.1016/j.compag.2004.04.003. URL http://www.sciencedirect.com/science/article/pii/S0168169904000705.

[64] Sharada P Mohanty, David P Hughes, and Marcel Salathé. Using deep

learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, 2016.

[65] Srdjan Sladojevic, Marko Arsenovic, Andras Anderla, Dubravko Culibrk, and Darko Stefanovic. Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience*, 2016.

[66] Lin Yuan, Jingcheng Zhang, Jinling Zhao, Shizhou Du, Wenjiang Huang, and Jihua Wang. Discrimination of yellow rust and powdery mildew in wheat at leaf level using spectral signatures. In *2012 First International Conference on Agro-Geoinformatics (Agro-Geoinformatics)*, pages 1–5, August 2012. doi: 10.1109/Agro-Geoinformatics.2012.6311599.

[67] Youwen Tian and Lin Zhang. Study on the Methods of Detecting Cucumber Downy Mildew Using Hyperspectral Imaging Technology. *Physics Procedia*, 33:743–750, 2012. ISSN 1875-3892. doi: 10.1016/j.phpro.2012.05.130. URL http://www.sciencedirect.com/science/article/pii/S1875389212014435.

[68] Matheus Kuska, Mirwaes Wahabzada, Marlene Leucker, Heinz-Wilhelm Dehne, Kristian Kersting, Erich-Christian Oerke, Ulrike Steiner, and Anne-Katrin Mahlein. Hyperspectral phenotyping on the microscopic scale: towards automated characterization of plant-pathogen interactions. *Plant methods*, 11(1):28, 2015. URL http://www.biomedcentral.com/content/pdf/s13007-015-0073-7.pdf.

[69] Sindhuja Sankaran, Reza Ehsani, Sharon A. Inch, and Randy C. Ploetz. Evaluation of Visible-Near Infrared Reflectance Spectra of Avocado Leaves as a Non-destructive Sensing Tool for Detection of Laurel Wilt. *Plant Disease*, 96(11):1683–1689, June 2012. ISSN 0191-2917. doi: 10.1094/PDIS-01-12-0030-RE.

[70] Fuan Tsai and William Philpot. Derivative Analysis of Hyperspectral Data. *Remote Sensing of Environment*, 66(1):41–51, October 1998. ISSN 0034-4257. doi: 10.1016/S0034-4257(98)00032-7. URL http://www.sciencedirect.com/science/article/pii/S0034425798000327.

[71] Ronald W. Schafer. What is a Savitzky-Golay filter?[lecture notes]. *Signal Processing Magazine, IEEE*, 28(4):111–117, 2011.

[72] J.-F. Huang and A. Apan. Detection of Sclerotinia rot disease on celery using hyperspectral data and partial least squares regression. *Journal of Spatial Science*, 51(2):129–142, December 2006. ISSN 1449-8596. doi: 10.1080/14498596.2006.9635087.

[73] Lin Yuan, Yanbo Huang, Rebecca W. Loraamm, Chenwei Nie, Jihua Wang, and Jingcheng Zhang. Spectral analysis of winter wheat leaves for detection and differentiation of diseases and insects. *Field Crops Research*, 156:199–207, February 2014. ISSN 0378-4290. doi: 10.1016/j.fcr.2013.11.012. URL http://www.sciencedirect.com/science/article/pii/S0378429013003997.

[74] Jing-Cheng Zhang, Rui-liang Pu, Ji-hua Wang, Wen-jiang Huang, Lin Yuan, and Ju-hua Luo. Detecting powdery mildew of winter wheat using leaf level hyperspectral measurements. *Computers and Electronics in Agriculture*, 85: 13–23, July 2012. ISSN 0168-1699. doi: 10.1016/j.compag.2012.03.006. URL http://www.sciencedirect.com/science/article/pii/S0168169912000804.

[75] Yingzi Du, Chein-I Chang, Hsuan Ren, Chein-Chi Chang, James O. Jensen, and Francis M. DAmico. New hyperspectral discrimination measure for spectral characterization. *Optical Engineering*, 43(8):1777–1786, 2004. ISSN 0091-3286.

[76] Jianwei Qin, Thomas F. Burks, Mark A. Ritenour, and W. Gordon Bonn. Detection of citrus canker using hyperspectral reflectance imaging with spectral information divergence. *Journal of Food Engineering*, 93(2):183–191, July 2009. ISSN 0260-8774. doi: 10.1016/j.jfoodeng.2009.01.014. URL http://www.sciencedirect.com/science/article/pii/S0260877409000259.

[77] Roberta H. Yuhas, Alexander FH Goetz, and Joe W. Boardman. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (SAM) algorithm. In *Summaries of the third annual JPL airborne geoscience workshop*, volume 1, pages 147–149. Pasadena, CA: JPL Publication, 1992. URL http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19940012238.pdf.

[78] E. Bauriegel, A. Giebel, M. Geyer, U. Schmidt, and W. B. Herppich. Early

detection of Fusarium infection in wheat using hyper-spectral imaging. *Computers and Electronics in Agriculture*, 75(2):304–312, February 2011. ISSN 0168-1699. doi: 10.1016/j.compag.2010.12.006.

[79] Anne-Katrin Mahlein, Ulrike Steiner, Christian Hillnhutter, Heinz-Wilhelm Dehne, and Erich-Christian Oerke. Hyperspectral imaging for small-scale analysis of symptoms caused by different sugar beet diseases. *Plant Methods*, 8(1):3, January 2012. ISSN 1746-4811. doi: 10.1186/1746-4811-8-3. URL `http://www.plantmethods.com/content/8/1/3/abstract`.

[80] T Rumpf, A-K Mahlein, U Steiner, E-C Oerke, H-W Dehne, and L Plümer. Early detection and classification of plant diseases with support vector machines based on hyperspectral reflectance. *Computers and Electronics in Agriculture*, 74(1):91–99, 2010.

[81] John Passioura. The drought environment: physical, biological and agricultural perspectives. *Journal of experimental botany*, 58(2):113–117, 2007. URL `http://jxb.oxfordjournals.org/content/58/2/113.short`.

[82] Christian Thurau, Kristian Kersting, and Christian Bauckhage. Yes We Can: Simplex Volume Maximization for Descriptive Web-scale Matrix Factorization. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, pages 1785–1788, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0099-5. doi: 10.1145/1871437.1871729. URL `http://doi.acm.org/10.1145/1871437.1871729`.

[83] Christoph Romer, Mirwaes Wahabzada, Agim Ballvora, Francisco Pinto, Micol Rossini, Cinzia Panigada, Jan Behmann, Jens Leon, Christian Thurau, Christian Bauckhage, et al. Early drought stress detection in cereals: simplex volume maximisation for hyperspectral image analysis. *Functional Plant Biology*, 39(11):878–890, 2012.

[84] Jan Behmann, Jorg Steinrucken, and Lutz Plumer. Detection of early plant stress responses in hyperspectral images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 93:98–111, July 2014. ISSN 0924-2716. doi: 10.1016/j.isprsjprs.2014.03.016.

[85] Dimitrios Moshou, Xanthoula-Eirini Pantazi, Dimitrios Kateris, and Ioannis Gravalos. Water stress detection based on optical multisensor fusion with a least squares support vector ma-

chine classifier. *Biosystems Engineering*, 117:15–22, January 2014. ISSN 1537-5110. doi: 10.1016/j.biosystemseng.2013.07.008. URL http://www.sciencedirect.com/science/article/pii/S1537511013001153.

[86] Kristian Kersting, Zhao Xu, Mirwaes Wahabzada, Christian Bauckhage, Christian Thurau, Christoph Roemer, Agim Ballvora, Uwe Rascher, Jens Leon, and Lutz Pluemer. Pre-symptomatic prediction of plant drought stress using dirichlet-aggregation regression on hyperspectral images. In *AAAI*, 2012. URL http://www-kd.iai.uni-bonn.de/pubattachments/620/kersting12aaai.pdf.

[87] Gerrit Polder, Gerie WAM van der Heijden, LC Paul Keizer, and Ian T Young. Calibration and characterisation of imaging spectrographs. *Journal of near Infrared Spectroscopy*, 11(3):193–210, 2003.

[88] Baolin Zhang and Douglas D Archbold. Water relations of a fragaria chiloensis and a f. virginiana selection during and after water deficit stress. *Journal of the American Society for Horticultural Science*, 118(2):274–279, 1993.

[89] Antonio Robles-Kelly and Cong Phuoc Huynh. *Imaging spectroscopy for scene analysis*. Springer Science & Business Media, 2012.

[90] Paul Geladi, Jim Burger, and Torbjörn Lestander. Hyperspectral imaging: calibration problems and solutions. *Chemometrics and intelligent laboratory systems*, 72(2):209–217, 2004.

[91] Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, et al. Fiji: an open-source platform for biological-image analysis. *Nature methods*, 9(7):676–682, 2012.

[92] Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1): 11–15, 1972.

[93] Daniel Cremers, Nir Sochen, and Christoph Schnorr. Towards recognition-based variational segmentation using shape priors and dynamic labeling. In *International Conference on Scale-Space Theories in Computer Vision*, pages 388–400. Springer, 2003.

[94] AK Mahlein, MT Kuska, S Thomas, D Bohnenkamp, E Alisaac, J Behmann, M Wahabzada, and K Kersting. Plant disease detection by hyperspectral imaging: from the lab to the field. *Advances in Animal Biosciences*, 8(2): 238–243, 2017.

[95] Puneet Mishra, Mohd Shahrimie Mohd Asaari, Ana Herrero-Langreo, Santosh Lohumi, Belén Diezma, and Paul Scheunders. Close range hyperspectral imaging of plants: A review. *Biosystems Engineering*, 164:49–67, 2017.

[96] Jan Behmann, Anne-Katrin Mahlein, Stefan Paulus, Jan Dupuis, Heiner Kuhlmann, Erich-Christian Oerke, and Lutz Plumer. Generation and application of hyperspectral 3d plant models: methods and challenges. *Machine Vision and Applications*, 27(5):611–624, 2016.

[97] Tawanda W Gara, Roshanak Darvishzadeh, Andrew K Skidmore, and Tiejun Wang. Impact of vertical canopy position on leaf spectral properties and traits across multiple species. *Remote sensing*, 10(2):346, 2018.

[98] L Ma and MK Upadhyaya. Effects of leaf position on reflectance, transmittance and absorption of red and far-red light in tomato, chenopodium album and amaranthus retroflexus leaves. *Weed Research*, 58(1):17–24, 2018.

[99] I Ibrahim. Illumination invariance and shadow compensation on hyperspectral images. 2014.

[100] Jonathon A Gibbs, Michael Pound, Andrew P French, Darren M Wells, Erik Murchie, and Tony Pridmore. Approaches to three-dimensional reconstruction of plant shoot topology and geometry. *Functional Plant Biology*, 44(1): 62–75, 2017.

[101] Aldo Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on pattern analysis and machine intelligence*, 16(2):150–162, 1994.

[102] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 1, pages 307–314. IEEE, 1999.

[103] Ke Lu, Xiangyang Wang, Zhi Wang, and Lei Wang. Binocular stereo vision based on opencv. 2011.

[104] H Christopher Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133, 1981.

[105] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.

[106] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.

[107] Jian Yao and Wai-Kuen Cham. Robust multi-view feature matching from multiple unordered views. *Pattern Recognition*, 40(11):3081–3099, 2007.

[108] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000.

[109] Rostam Affendi Hamzah, Rosman Abd Rahim, and Zarina Mohd Noh. Sum of absolute differences algorithm in stereo correspondence problem for stereo matching in computer vision application. In *Computer Science and Information Technology (ICCSIT), 2010 3rd IEEE International Conference on*, volume 1, pages 652–657. IEEE, 2010.

[110] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2010.

[111] Changchang Wu. Towards linear-time incremental structure from motion. In *3D Vision-3DV 2013, 2013 International conference on*, pages 127–134. IEEE, 2013.

[112] Changchang Wu, Sameer Agarwal, Brian Curless, and Steven M Seitz. Multicore bundle adjustment. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3057–3064. IEEE, 2011.

[113] Michael P Pound, Andrew P French, Erik H Murchie, and Tony P Pridmore. Surface reconstruction of plant shoots from multiple views. In *European Conference on Computer Vision*, pages 158–173. Springer, 2014.

[114] Josien PW Pluim, JB Antoine Maintz, and Max A Viergever. Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, 22(8):986–1004, 2003.

[115] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In Vittorio Scarano, Rosario De Chiara, and Ugo Erra, editors, *Eurographics Italian Chapter Conference*. The Eurographics Association, 2008. ISBN 978-3-905673-68-5. doi: 10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136.

[116] Emanuel Parzen. On estimation of a probability density function and mode. *Ann. Math. Statist.*, 33(3):1065–1076, 09 1962. doi: 10.1214/aoms/1177704472. URL `https://doi.org/10.1214/aoms/1177704472`.

[117] Abraham Savitzky and Marcel JE Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8): 1627–1639, 1964.

[118] V. S. Manivasagam and R. Nagarajan. Water stress assessment in maize crop using field hyperspectral data. 2015.

[119] Anne-Katrin Mahlein, Ulrike Steiner, Christian Hillnhütter, Heinz-Wilhelm Dehne, and Erich-Christian Oerke. Hyperspectral imaging for small-scale analysis of symptoms caused by different sugar beet diseases. *Plant methods*, 8(1):3, 2012.

[120] Faxian Cao, Zhijing Yang, Jinchang Ren, Mengying Jiang, and Wing-Kuen Ling. Does normalization methods play a role for hyperspectral image classification? *arXiv preprint arXiv:1710.02939*, 2017.

[121] AW Johnson, MJ Davies, OM Grant, and DW Simpson. Developing strawberry cultivars with improved water use efficiency. *Acta Hortic*, 1049:853–858, 2014.