An investigation into apple inspection in colour space

A thesis presented in fulfilment of the requirements for the degree of

Master of Engineering

in

Industrial Automation

at Massey University,
Palmerston North, New Zealand.

Michael Caulton

2011

Abstract

To maximize the storage life of packed apples, no damaged apple can be included because the ethylene produced by the damaged, rotting fruit can lead to damage in the surrounding fruit. The variability in size, shape and colour of apples has meant that the commercial inspection process has had to remain largely manual. In manual inspection, all surfaces of the apple are sequentially presented to human inspectors who remove deformed or blemished apples from the processing line.

Most automated commercial inspection systems use infrared images to inspect the apples. These systems struggle to identify defects near the stem or calyx because of the reduced information in the monochrome images. In this work, the full visible light spectrum, via a full colour camera, was used to inspect apples in an attempt to isolate and identify stems, calyxes and defects.

A Sobel filter and an islanding routine were used to find areas of interest in undamaged apple skin. This led to the identification of 92.7% of stems, 97.9% calyxes and 83.0% of defects. Image processing for defect identification took as little as 16ms per image; fast enough for implementation in a commercial application where speed is critical.

Acknowledgments

I would like to thank my supervisor Dr Rory Flemmer for his guidance and willingness to share his knowledge throughout the course of my research.

I would like to express my gratitude for the time spent by my co-supervisor Dr Claire Flemmer, critiquing and offering advice on my thesis. Your perseverance was greatly appreciated.

The research was financially sponsored by Steve Saunders of GroPlus Ltd, Tauranga. Thank you for the financial help over the period of the research.

Table of contents

1. Introduction

55 Million tons of apples are grown annually, worldwide (O'Rourke, 1994). Whether or not they spend time in a controlled-atmosphere storage facility, the vast majority of them are packed into containers, mostly pre-formed trays. This process is partially automated in that apples are floated out of the bin, washed, waxed and then fed to the main conveyor which transports the apples down the length of the packing shed. In most packing sheds, the main conveyor is intelligent in that it classifies apples by weight and by colour and then has the capability to discharge a particular apple at a particular point as it moves down the length of the shed. At each such discharge point, and there may be up to fifty of them, there is a transverse “tray conveyor”. This passes orthogonally under the main conveyor and workers load trays onto this conveyor on the far side of the main conveyor. At each “drop” one or several workers take apples and pack them into trays. Often the drops are so arranged that the apples roll more or less into the pockets of the tray under gravity and then workers merely tidy up the process.

For all this automation, the industry is currently very heavily reliant on labour to pack into trays. The other area which is labour intensive is the inspection of the fruit. This occurs just before the apples are introduced onto the main conveyor. The apples advance along a set of rollers perhaps 1200mm long, over a distance of perhaps 2000mm. The rollers rotate as they advance. In this way, all surfaces of the apple are presented to the eyes of human inspectors who remove deformed or blemished apples. In the nature of human inspection, extended as it is over several hours at a time, performance is not perfect and is variable depending on inspectors and how long they have been on duty. This is problematic for the installation of automated packing from the drops because there is no longer the human backup to find bad

apples before they are put in the box. Because of ethylene production by bad apples, the proverbial bad apple can damage its neighbours.

As discussed at length in the literature review, automated inspections systems are on the market. Unfortunately, these are only capable of pulling out the gross defects. This makes the inspection process easier because the inspectors are pulling out fewer apples and therefore spend less time being distracted from their inspection task. Despite this automation, packing shed operators are aware that inspection is not adequate. Further, in the nature of automation, it is mandatory that each stage of the process produces an acceptable output; it is not wise to have to have a safety net to deal with an inadequate upstream process.

Commercial inspection systems use many infrared images of the apple to do the inspection. By this process, they are working with a monochromatic image which is not sensitive to colouring of the apple. These images can be thresholded and all dark pixels declared to be potentially unsavoury. The difficulty arises that there is always a dark area around the stem and calyx, which might be and usually is benign. Systems use geometrical information from the many images to find two dark areas approximately opposed diametrically and declare these to be either stem or calyx – areas of prescribed radius in which no inspection can occur. They then declare any other area to be a blemish and threshold on the basis of area. This system has two fatal defects. Firstly, there are many defects which occur particularly around the stem but also in the calyx. These cannot be seen. Secondly, blemishes are thresholded on their size, which is to say, visual impact, alone. But there are many defects, such as stem punctures which are small but lethal to the apple.

The need for automated inspection therefore arises in the first instance from the desire to automate the process of packing apples into trays and this process demands an adequate upstream inspection process. In the second instance, current technology, relying on thresholding of infrared images is fatally compromised because it cannot inspect the area

around the stem and calyx, nor can it inspect and classify defects on their possible consequences. These inadequacies lead pack house operators to have a safety net of human inspectors but, unfortunately, the mesh of this net is not fine enough.

These considerations have led to this broad investigation of whether it is possible to use colour images to rectify the shortcomings of current automated inspection. These preliminary results indicate that it is possible and that the technology can be commercially deployed as many research results reported in the literature cannot.

2. Literature Review

2.1. Introduction

Customers judge the quality of fresh fruit and vegetables primarily on their external visual attributes (Kader, 1992). Trays of produce of uniform size, shape, colour and orientation are perceived, by the consumer, to be of a higher quality than trays that are non uniform. Further, produce with visual defects on the surface are perceived to be poorer quality than blemish-free produce (Blasco, Aleixos et al. 2007). Consequently, an important step in the post-harvest process is inspection and grading of produce.

Inspection can be used to measure attributes such as size, density, colour and presence of surface defects as well as internal attributes such as sugar content and dry matter. Inspection techniques are either invasive (for example, a penetrometer test to judge fruit firmness) or non-invasive (for example, vision-based inspection by humans or by automated systems), the latter having the advantage that the inspected produce is not damaged and can therefore be on-sold. Human inspectors are subjective, susceptible to fatigue and variable. This review focuses on non-invasive automated inspection techniques. Of these, only computer vision-based inspection systems are commercially viable and research in this area is reviewed in detail, together with the work which has been done on presenting the entire external surface for inspection.

2.2. Summary of non-invasive inspection technology

Non-invasive inspection technologies are reviewed by Studman (2001) and by Van Zeebroeck et al., (2007). They include the following:

- X-ray images to find internal defects such as rind-puffing in oranges (Njoroge et al., 2002). The emitted X-rays pass through the fruit and onto an X-ray scintillator. This

device has a thin layer of luminescent material which allows a monochrome camera to capture the X-ray image for processing. The processing time and effectiveness of the method is not reported.

- Nuclear Magnetic Resonance (NMR) to identify the pits in olives as they move along the packing conveyor (Zion et al. 1997). As the olives move along a conveyor made of non magnetic materials, the system acquires NMR projections with a slice thickness of 3mm using a superconductor magnet and an imaging coil. In a test of 300 olives, half of which were pitted, the misclassification error was between 2.3% and 4.7% for a conveyor speed of 0mm/second (stationary) to 250mm/second. The system could process approximately 10 olives per second which is very low for an industrial system. This throughput could be increased through the use of a longer imaging coil. Unpitted olives are currently easily separated using flotation which reduces the number of unpitted olives reaching the packing line to below 1%. This percentage is further reduced by manual inspection. It should be noted that the NMR system is prohibitively expensive and has a misclassification error of 4 times the expected error of the flotation system.
- Proton Magnetic Resonance (PMR) to detect internal browning in apples (Chayaprasert and Stroshine, 2005). The apples travel on a non magnetic conveyor belt through a Plexiglas tube wrapped with a magnetic resonance probe coil. The tube is surrounded by a permanent magnet with field strength of 0.13T. Tests on a small sample of apples showed that the PMR signals of healthy apples decayed slower than that of apples with internal browning. Both the speed of the conveyor and orientation of the fruit influenced the results. A conveyor speed of less than 50mm/second had the lowest classification error of about 12% while errors were as high as 76% at a conveyor speed of 250mm/second. The tests concluded that if the apple orientation and a conveyor speed of less than 100mm/second were accurately maintained the

PMR system could, with reasonable accuracy, identify apples with internal browning. Unfortunately the conveyor speed is far too slow for an industrial application and constant apple orientation is very difficult to achieve in industrial pack houses.

- A Spectrophotometer with a waveband of 400nm-1700nm (thus covering both NIR and visible light) was used to detect fresh bruises, less than 45 minutes old, in apples (Xing and De Baerdemaeker, 2007). A sample of 178 Golden Delicious, 95 Jonagold and 56 Braeburn apples were used. The apples were impacted using a pendulum with an impact energy of approximately 0.11J. Multispectral images were taken of the affected area and surrounding sound flesh for later analysis. The fruit's E-modulus (a measurement of stiffness), maximum force and bioyield force were then measured using a universal test machine (UTS Tesesystemse GmbH, Ulm, Germany) with an 8mm cylindrical plunger. Both bruised and sound flesh was measured. The measurements showed sound flesh always had a higher E-modulus than bruised flesh. The UTS measurement is destructive which limits its use in a commercial application. The authors tried to find a correlation between the E-modulus and the spectral reflectance. No dominant wavebands were identified but light in the visible spectrum could not be used on bi-coloured apples and reflectance was too low for light with wavelengths below 744nm and above 1400nm. The use of the calculated E-modulus as a classification method had an accuracy of more than 95% across the three varieties.
- Computer vision to measure a wide range of parameters including size, external colour and blemishes for a variety of produce. This is reviewed in more detail in sections 2.3 and 2.4.

While many of these types of measurements ‘work’, in the sense that they do yield accurate inspection, they are not currently commercially viable because they are too expensive and

too slow. The only commercially significant measurement system is computer vision-based inspection technology which can measure size, colour and a limited range of defects if they occur outside the area around the stem and the calyx.

2.3. Computer vision-based inspection research

Research into computer vision-based inspection of produce began more than twenty years ago and has led to commercial vision inspection systems, such as the InVision 5000 by Compac, New Zealand[1] and the ISHII Apollo Grader and Sizer (Njoroge et al., 2002), becoming a common feature in pack houses.

Most non invasive inspections systems use a computer to process data captured from cameras. The processing speed has to be able to handle the high throughput of the commercial facility and this often results in low quality images being used in the inspection systems. The lower the quality of the image the less likely any defects are to be identified (Yang and Marefat 1994). With the advent of multi core computer processors this has becomes less problematic but the image quality and processing algorithms are still significantly limited by computer architecture.

Camera based technology allows external features, such as colour, size and an assortment of defects, to be identified. Several fruit have very unique shapes; a pear is tapered towards the stem and rounded at the bottom whereas an orange is nearly perfectly spherical. An apple is described as a 'tilted elliptic spindle torus' (Narayanan et al., 2008) with a concavity where the stem and calyx are located. The range of fruit shapes means that computer based sizing, inspection for defects and grading can be difficult. In order for an apple to be

[1] (Compac. 2010. Sorting, grading and sizing fruit and vegetables, fruit graders, colour sorters[Online] Available: http://www.compacsort.com [accessed 15 January 2011])

accurately inspected with computer vision, attributes such as the stem and calyx must to be correctly identified so that they can be eliminated as possible defects.

Inspection for defects relies upon capturing the complete surface image of the produce. The problem of exposing the entire external surface of the fruit for inspection and of identifying and isolating the stem/calyx areas from the inspection are discussed in section 2.4.

The research work is reviewed chronologically below.

2.4. Inspection technology

Davenel et al. (1988) describes a vision system designed to identify defects in Golden Delicious apples. The system used a conveyor of bi-cone rollers to tumble apples under a 6-bit monochrome camera with a resolution of 208 x 144 pixels and used a 100nm bandwidth filter centred on 550nm. The processor speed available at the time significantly limited the performance of the system. In order to achieve the desired throughput a processor speed of 900MHz was required while only an 8MHz processor was available. To overcome this problem the image resolution was reduced to 52 x 52 pixels. The automatic grading system correctly graded 69% of fruit with 26% of fruit graded just above or below the correct grade. This was considerably lower than European standards on acceptable grading error, so the system could only be used on low grade apples. Fruit was oriented mechanically prior to inspection.

In order to detect bruises in Golden Delicious apples a spectrophotometer was used by Geoola et al. (1994). The spectrophotometer had a range of source wavelengths from 400nm to 840nm with a resolution of 2nm. A sample of 300 apples that had been in storage for varying lengths of time, up to six months, was used. The sample was split into three different batches, namely, 100 unbruised stored apples, 100 stored apples that were bruised and maintained at room temperature for 24 hours and 100 stored apples that were bruised and

maintained at room temperature for 90 minutes. The bruises were manually inflicted using a steel disk, 39mm in diameter and weighing 364g, to impact the fruit from heights of 10, 15 or 20cm. Source wavelengths ranging from 750nm to 800nm produced the best results for bruise detection in Golden Delicious apples. The system had an accuracy of 96.1%, 88.4% and 93.7% for the three different apple batches respectively. The same experiment was then performed on a new sample of apples that were fresh and unrefrigerated. This had an improved accuracy of 92.9% for the unbruised apples and 91.8% for 24hour old bruises.

The use of a colour camera to identify defects in Golden Delicious apples was proposed by Leemans et al. (1998). The system used a low resolution, 240 x 180 pixel, colour camera to apply three filters to an image of an apple to identify defects. The sample library consisted of 80 apples specifically orientated so the stems and calyxes were out of view, perpendicular to the camera's optical axis. The initial filter compared each pixel on the apple to a median normal colour using the Mahalanobis distance between the pixels. If too great the pixel was identified as defective. The second filter grouped the defect pixels according to their Mahalanobis distance into their closest class of which the weighted distance was the smallest creating global averages. This helped identify diffuse defects but also increased the portion of healthy skin identified as defect. A third filter was introduced to reduce this miss-classification using a localised segmentation routine.

Multi spectral inspection systems have been developed to try and take advantage of the benefits of more than one light spectrum. For example a hermetic box was created allowing identical images to be captured in parallel, for computer vision analysis, in both colour and infrared spectrums (Aleixos et al., 2002). The hermetic box used a splitter prism and mirror to split the input light so both a polarized colour camera and an infrared filtered mono chrome camera could capture the same image. Each camera was run on a separate digital signal processing (DSP) board running in parallel with one another. This allowed the colour and

infrared inspection routines to be processed in parallel, reducing the overall processing total time per image. This hermetic box in conjunction with a fruit tumbler conveyor belt and lighting apparatus was used to determine the size, shape, colour and presence of defects in citrus fruit. Both mandarins and lemons were used for testing of the system. The system was designed to be used in industry but did not meet the minimum speed requirement of 10 fruit per second with the average processing time of 124ms per fruit. It was proposed that with minor hardware upgrades the 10 fruit per second rate could be achieved if the fruit being processed was medium sized and had minimal defects while 16 fruit per second could be achieved if the fruit were smaller and had minimal defects.

Li et al. (2002) and Blasco et al. (2003) describe an inspection system for fruit (oranges, peaches and apples) held in a cup which is rotated under a camera, allowing for most of the fruit's surface to be captured in as little as 4 images. The system used a 768x576 pixel colour camera and a ring shaped fluorescent tube with dome diffuser. The camera was positioned so it had a relatively low resolution of 3.5mm per pixel. The process takes non-overlapped images which allows for much of the surface to be inspected quickly. Colour segmentation was used in the image processing to estimate the size of the fruit and identify defects on the skin. The stem detection and location algorithm had an error of 7.5% for oranges, 8.6% for peaches and 5.3% for apples. During online testing the system identified 86% of external defects correctly. The downside to this mechanical arrangement is that it can only process a maximum of 3-4 fruit per second, whereas industry would like throughputs of at least 10 fruit per second. In addition, part of the fruit is occluded by the suction cup and is consequently not inspected.

High spatial resolution hyperspectral imaging was used by Mehl et al. (2004), to identify defects and contaminations on Red Delicious, Golden Delicious, Gala and Fuji apples. They used a monochrome CCD camera equipped with a spectrograph to capture sample images.

Three different methods were used to analyze the images; the second difference method, the second difference – asymmetric model, and the direct monochromic method. The second difference method was used to identify the presence of weak and strong absorption peaks of the spectrum. The second difference – asymmetric method used a generalized asymmetric formula to identify the difference between the intensity signal and estimate signal intensity. The monochromic comparison analyzed the images in three spectral bands; 542, 682 and 752nm. The work concluded no single band wavelength could easily be used to identify defects in the monochromic images and the red chlorophyll absorption band on its own was insufficient for separating defects or contaminations from normal skin. The asymmetric difference method was the most accurate at identifying defects, particularly bruises. The method was independent of apple variety and fast to process as it only used a single wavelength.

Most machines used to grade and sort apples use mechanical bi-cone rollers to tumble the apple under a camera. Throop et al. (2005) felt that this was not the best method to analyze an apple because of the difficulty in locating the stem and calyx. They designed a conveyor system that orientated the apple with the stem and calyx in a constant position, with 97.6% accuracy. A camera was then positioned so the stem and calyx region of the apple would not be inspected. The camera used several linear filters mounted in front of a camera lens to provide images for three different wavebands (740nm, 950nm and visible). LED lighting was arranged asymmetrically to correct for the variation in geometric reflectance due to the spherical nature of an apple. LED lights have no noticeable decay and a greater life span (20 000 hours) compared with noticeable decay and 750 hour life of incandescent lights. The system had a grading error of 12% over a sample of 959 apples. This is outside the 10% error accepted by the USDA (United States Department of Agriculture) grading standard but with improvements the system was expected to achieve fewer than 5% misclassification.

Many varieties of apples are bi-coloured and this makes the image processing component of a machine vision system more challenging. Kleynen et al. (2005) used a high resolution (1280 x 1024 pixel) monochrome camera and four IR filters centred about 450nm, 500nm, 750nm and 800nm, to inspect bi-coloured Jonagold apples. They created a template of an average stem and an average calyx and applied segmentation and template matching to an image library with 280 sound apples, 246 defective apples and 292 stem and calyx views. Stems and calyxes were identified to an accuracy of 91% and 92% respectively using the 800nm filter. Their method had a success rate of between 94.5% for seriously damaged fruit down to 55% for fruit with slight damage. The study concluded that the 500nm wavelength offered no useful information while the 450nm band was useful at identifying slight surface defects and the 750nm and 800nm bands offered good contrast between other defects and sound apple skin.

A machine vision system was developed by Bennedsen and Peterson (2005) using NIR and monochrome images to identify defects on pre-oriented apples of eight different varieties. The apples were oriented with the stem and calyx perpendicular to the cameras. Two separate vision systems were used; the first confirmed the orientation of the apples while the second did the defect inspection. The apple orientation detection was performed using Matlab's neural networks toolbox (The math works, Inc. Natick, MA, USA). The images had a contrast filter applied to reduce the contrast and reduce small variations in grey level. For defect detection an optical splitter was placed in front of the camera so three identical images could be captured simultaneously through the different filters. One image remained unfiltered while the other two had a 740nm cut filter and 950nm cut filter respectively. The 740nm filter identified various diseases more effectively while the 950nm filter identified bruises well. The majority of the sample fruit had bruises inflicted the day prior by dropping them from a height of 150-200mm onto a convex surface. Four different segmentation routines were used to

identify defects in the sample fruit. The orientation verification vision system had an accuracy of greater than 99% while the defect detection vision system had an accuracy of between 77% and 91.6% over the eight varieties of apples tested. This may be adequate for industry depending on the existing quality criteria.

Unay and Gosselin (2006) used supervised segmentation to identify apple defects in bi-coloured apples and tested their method on the same image database used by Kleynan et al. (2005). Image segmentation used eight neighbourhood tests ranging from one to nine pixels in distance from the point of interest. The segmentation algorithm became too computationally expensive for large numbers of pixels of interest, so thresholding was used in conjunction with classifiers to segment the images. The two supervised classifiers that worked best were the Support Vector Machines (SVM) and Multi-Layer Perceptions (MLP). The computation time for the 430x560 pixel image was 280ms, 341ms and 1,081ms respectively. The project aimed to process 10 fruit per second and estimated that three images of each fruit would be required to capture the complete apple surface. This implies a required computational time of 33ms per image. By reducing the image resolution to 53x71 pixels, the MLP computation time dropped to 51ms per image. However, with the reduced image resolution, the error in defect detection increased. The average error across all defect types was approximately 15%.

Similar segmentation techniques were applied to inspection of citrus fruit by Blasco et al. (2007b). This work used unsupervised segmentation methods which do not need training over the course of the season to account for variation in colour. A colour camera, with a resolution of 768x576 pixels was used to capture images of 356 oranges and 279 mandarins. Cross polarization was used to reduce bright spots in the images. Pre-processing of the images to reduce them from 8 bit colour images to 5 bit colour images improved the computational speed of the algorithm without degrading the results. An adaptation of the

Generalized Lloyd Algorithm (Lloyd, 1983) was used to perform the colour reduction. The algorithm was tested against the 2, 294 images of defects and stems with a success rate of 94%. Unfortunately the computational time is not reported so a direct comparison with the previous work (Unay and Gosselin, 2006) is not possible. It is however likely that the improvement in inspection capability arises from images having over 100 times the resolution.

When the skin of a citrus fruit is damaged it releases an essential oil. These oils are fluorescent and can easily be seen under ultra violet lighting conditions. Typically pack houses use special rooms illuminated with ultra violet light and operators must wear protective glasses and gloves because prolonged exposure to ultra violet radiation is dangerous to human skin. The system proposed by Blasco et al. (2007a), used multispectral computer vision to eliminate this health hazard and reduce the number of people on the processing line. The system used three different cameras to capture images in visible colour, near infrared and ultra violet spectrums. Camera resolution ranged from 768x576 pixels to 756x581 pixels. In the visible colour range, only 63.4% of defects were successfully identified in a sample size of 336 fruit, in the ultraviolet range 79.5% of defects were identified in a sample size of 39 fruit and in the near infrared range 92.9% were identified in the sample of size 56 fruit. Of the six different defects analyzed, only the stem end injury could be identified by the algorithm because it was the only one that was visible under ultra violet light. The algorithm could not distinguish between the other five types of defects.

Many horticultural products are spherical in nature. This results in images that tend to be darker near the edges than in the centre. This occurs because the distribution of light radiation being reflected from the produce and absorbed by the cameras sensor is heavily dependent on the position of the light source and the geometry of the produce. It was shown that monochrome hyper spectral images of citrus fruit, mandarins in particular, could have a correction factor applied to minimize this effect (Gomez-Sanchis et al., 2008). The

methodology considers the fruit to be a Lambertian ellipsoidal surface in order to calculate the part of radiation that should reach the camera's sensor. This allows for the radiation intensity for all pixels to be uniform across the complete spherical surface thus allowing larger portions of an image to be accurately processed in computer vision applications.

Lee et al. (2008) report a commercially viable automated date grading system which uses two near infrared (750-1200nm) filtered cameras to process four lanes of dates. The dates move on a conveyor at a speed of 406mm/second and sharp images are taken with the camera shutter speed set at 8 milliseconds. Eighteen 10W diffused halogen lights illuminate the four lanes. The system grades dates into four different categories dependent on size and the amount of skin delaminating. Less skin delaminating is more desirable. The average human grading accuracy is between 60% and 72%. The automated vision system had an accuracy of 74% to 79%, a gain of approximately 10%.

Pomegranate is a fruit that has excellent nutritional properties but is difficult to peel which affects the marketability and decreases its potential for consumption. Ready to eat pomegranate arils is one way to increase the consumption of the fruit but if the internal membrane or defective arils are mixed with good arils the produce's shelf life and perceived quality deteriorates. A machine designed to sort the desired arils from undesired was proposed by Blasco et al. (2009b). The system used four conveyors monitored by two colour cameras and a computer to sort the pomegranate. A blast of air, aimed at a collection chute, was used to remove a particular object from the conveyor. Two different segmentation algorithms were used; simple thresholding and Bayesian linear discriminating analysis. The conveyor belts were blue while the pomegranate internal membranes are white and the pomegranate arils vary between white and red. These contrasting colours meant the simple R/G ratio thresholding worked well with a 90% success rate while being less computationally

intensive than the Bayesian linear discriminating analysis. The system was implemented in industry and had a maximum throughput of 75kg/hour.

Processes such as the canning of mandarin segments are extensively automated. The fruit is peeled, the segments separated and each segment's skins removed. This is all performed by machines but each segments must then be manually inspected and sorted so as to ensure any remnants of peel or skin membrane, broken segments, non separated segments and segment with pips are removed from the processing line. A mechanical system was designed that used semi transparent singulator conveyors illuminated from below to allow a computer vision system to inspect and sort these segments (Blasco et al., 2009a). The system was designed for commercial use, running at 50 fruit segments/second. The computer vision component of the system used a single core computer with two colour cameras acquiring the images at a relatively low resolution (512x384 pixels). The two primary discrimination algorithms used to sort the segments were a segment symmetry overlay and a radius contour signature plot. The segment symmetry overlay split the segment about its longest axis and compared the two halves to ensure the segment was symmetrical and the end of the segment was not broken off. The second method recorded the radius of the segment at equal intervals resulting in a radius signature that could be used to identify damaged segments. Broken segments are common, making up about 18% of the segments processed. The system, when implemented on an industrial packing line, correctly identified 96% of segments that contained pips or pieces of skin, 93% of segments that were sound and 83% of broken segments.

2.5. Fruit orientation

Fruit inspection technology frequently fails because the stem and calyx are mistaken for blemishes. Several researchers have studied ways to orient fruit in order to avoid this problem.

Penman (2001) demonstrated that the reflected light pattern, captured by an overhead camera, from a long narrow linear blue light source, shone at a tumbling apple could be used to detect the concavity of the stem or calyx. A blue light source was used because apples typically do not have a blue colour component. This allows for the concavity detection to be used in conjunction with colour based blemish detection. The method does not identify what the concavity is (stem, calyx or defect) but identifies the region where a focused search for defects can be initiated. The centre of the stem or calyx concavity was typically located to within 4mm when the stem/calyx concavity was near the centre of the camera's view. The further the stem/calyx concavity moved away from the centre of the camera view, the less accurate the detection became.

It is considered to be more desirable to have a single camera for inspection because of the reduced cost and complexity of the system. For this reason several studies into the feasibility of methods to capture a complete surface image of an apple, through a single image of the fruit, have been performed. The use of parabolic mirrors oriented in several different arrangements allowed for a complete surface image of an apple to be captured in a single image (Reese et al., 2009). The reflected images of the apple were captured from the mirrors via a camera mounted directly above the apple. The use of parabolic mirrors as opposed to flat mirrors increased the resolution while decreasing the distortion at the edges of the images. The drawback of mirrors being used in this industry is that they must be kept clean. This is important so as not to introduce artefacts into the captured images being analyzed. Although the food industry requires high levels of cleanliness, clean room status is not feasible and there would generally be enough dust present on the mirrors to create artefacts. For this study the apples were manually rotated while being suspended on two taut parallel wires in order to minimize the area of apple obstructed by any structural members. This method is not commercially viable as it can damage the fruit and there is very little control

over the fruit while suspended on these wires. Reese et al. (2009) note that the wires could be replaced with larger parallel rails resulting in only a small loss of surface area. Alternatively the top half of the apple could be processed, the apple rotated on the rails and then the bottom half could be processed.

Pre-orientated apples can be presented to the cameras so the position of the stem and calyx are known and thus not misidentified as defects. Because of an apple's shape and mass distribution its internal properties can be used to orientate the apple so the axis of the stem and calyx is known. Apples allowed to roll, propelled under the force of gravity, down a track consisting of two parallel rails were shown, theoretically and experimentally, to orientate in a manner where the stem and calyx was parallel to the rails and perpendicular to the direction of the travel (Lefcourt et al., 2009). The track consisted of two parallel rails 1.2m long and on an incline of 12 degrees. A monochrome camera running at 60 frames per second was used to record the behaviour of the apples as they rolled down the track. The weight, diameter, height, smallest and largest radius of 200 sample apples was recorded. The apples were then placed on the tracks using one of two methods. In the first method the apples were randomly orientated and poured out of a container onto the rails. The second method used an existing plastic cup-type chain driven conveyor to deposit apples onto the rails. The method had a first pass success rate of 80%, i.e., 80% of the apples were correctly oriented by the time they reached the end of the rails.

3. Materials and Methods

3.1. Experimental Apparatus

An experimental rig was built in order to capture a library of images of apples. The rig is shown below in Figure 3.1.

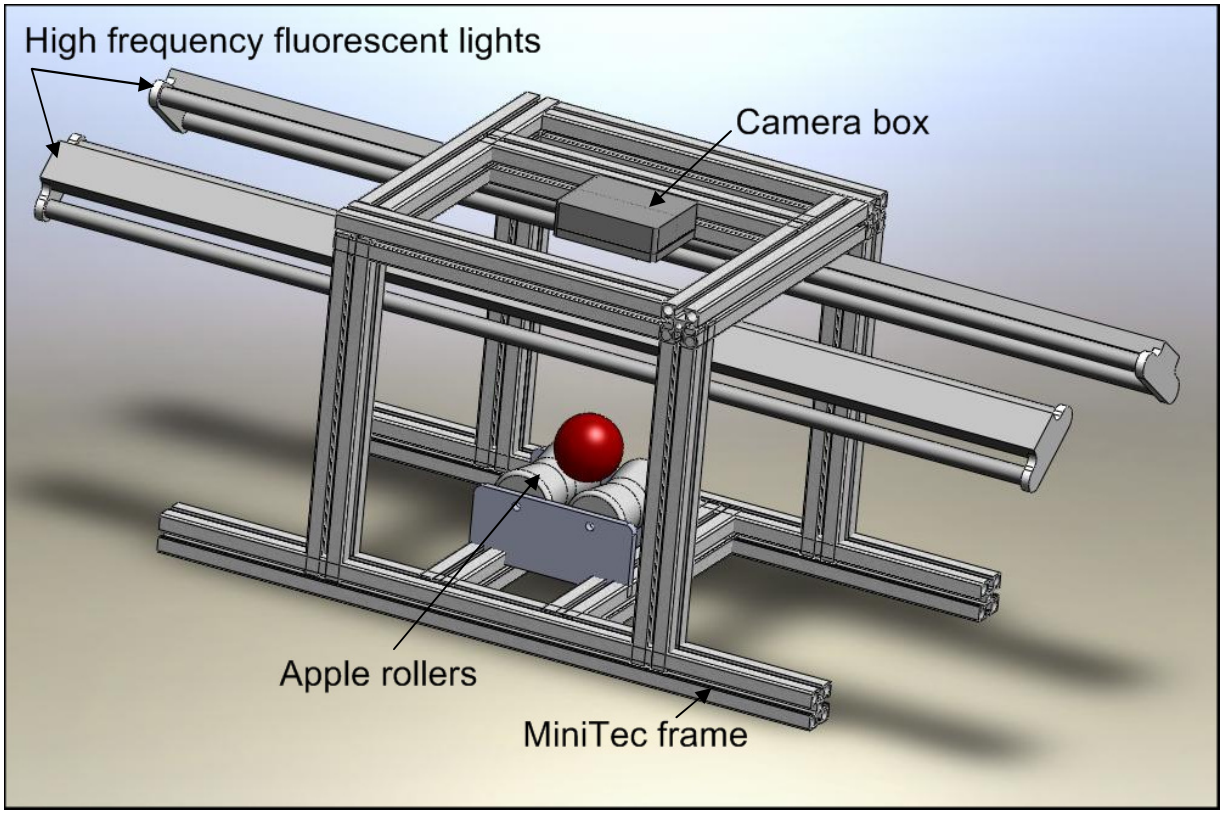

Figure 3.1 Diagram of experimental rig

The mechanical structure supporting the cameras, lights and rollers was constructed from MiniTec 20.1033 45x45F aluminium extrusion. This type of structure allows quick and easy adjustment of the component positions to find the best configuration. The structural components near the rollers were painted matt black to minimize reflected light to the cameras.

3.1.1. Light Source

For machine vision applications the type of light source is limited because light flicker is a problem when capturing an image. Standard fluorescent lamps with magnetic ballasts flicker at twice the supply frequency. In New Zealand the supply frequency is 50 Hz so the fluorescent will flicker at 100 Hz. The Logitech camera has an effective shutter speed of a sixteenth of a second so the brightness as observed by a random sample of images will vary considerably.

As a consequence of this, a high frequency (HF) fluorescent lamp ballast was used. These flicker at about 20 kHz or 1300 times an image which negates the effect they might have on the brightness of an image. Two 6 foot long dual HF fluorescent lamps were used with Philips TLD 36W/865 Cool Daylight tubes. The lamps were placed at an angle of 60 degrees from the horizontal axis to maximize the light hitting the apples surface.

Fruit such as apples have very reflective skin due to wax on the surface (it is standard practice to wax apples to reduce desiccation over storage life). Light shone on the skin is reflected and causes bright spots, called ‘specular reflections’, in the image. In order to remove this effect, a linear polarizing (LPL) filter was placed in front of the lamp to cross polarize the light (Figure 3.2). This arrangement ensures that only absorbed light passes through to the camera’s sensor. The filter used to polarize the light was a linear Polyvinyl Alcohol-Iodine filter. To remove the effect of external light sources, all images were captured in a dark room with illumination provided solely from the fluorescent lamps.

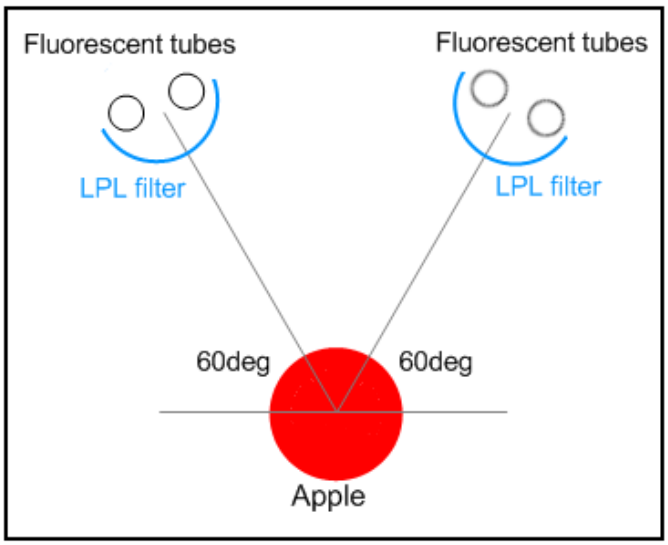

Figure 3.2 Lighting arrangement for removal of specular reflections

3.1.2. Cameras

A Logitech Quickcam 9000 Pro webcam was chosen for image capture because it is inexpensive, provides adequate image quality and has an acceptable noise level. The camera uses a CMOS (complimentary metal-oxide semiconductor) sensor with a resolution of 960x720 pixels. It has a built in Carl Zeiss lens with a short focal length (3.7mm), a lens iris F factor of 2.0 and automatic focus adjustment. A high speed USB2.0 connection linked the camera to a computer and allowed for 15 frames per second (FPS) to be captured at full video resolution. The video stream from the camera was captured using the Windows DirectShow API (Application Programming Interface) in conjunction with Geraint Davies' FSFWrap.dll and Capstill.dll. (Davies, 2005) This allowed for individual frames to be captured from the video stream and analyzed inside Microsoft Visual Basic 6.

Two of the cameras were mounted inside enclosures fitted with light filters (Figure 3). A Marumi 28mm Circular Polarizing (CPL) filter was used to remove specular reflections (from the shiny apple skin) in the captured colour images. A 28mm 850nm Infrared (IR) pass filter

was used to capture monochrome images in the Infrared spectrum, the Logitech camera's internal infrared block filter having been removed.

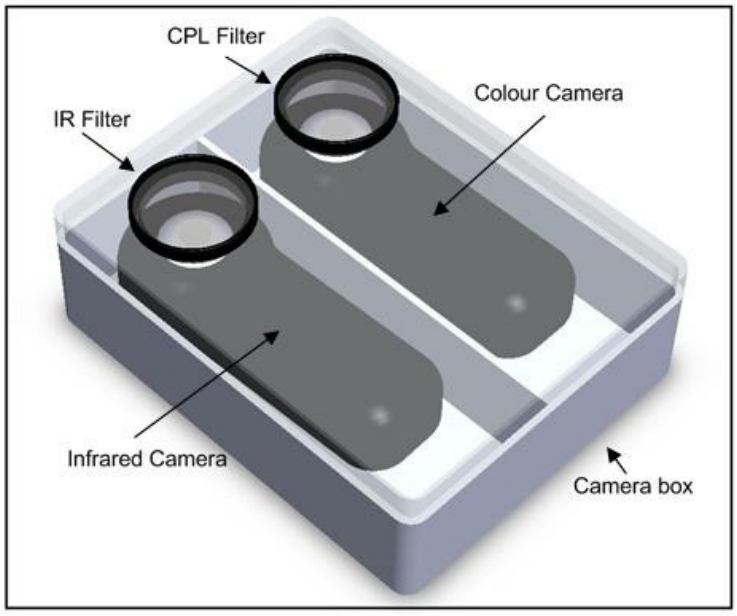

Figure 3.3 Camera enclosure model with lid made transparent

3.1.3. Computer specifications

The computer used to perform all image processing was an Apple iMac running Microsoft Windows XP Professional service pack 3. The iMac had an Intel® Core™ 2 Duo CPU E8235 at 2.80 GHz and 2.98 GB of RAM. Microsoft Visual Basic 6.0 for 32-bit Windows was used for the software development.

3.1.4. Fruit handling and camera position

In designing an experimental rig for capturing apple images, the aim was to choose a design which would be viable in an industrial pack house where apples tumble as they move along a conveyor. This places constraints on camera position; cameras cannot be placed too near the horizontal, otherwise their view will be obscured by the conveyor rollers. Similarly, the angle of view of the camera must be limited otherwise images will show apples in

neighbouring conveyors. With these constraints, the apple was placed on top of two rotating rollers with camera arrangement shown in Figure 3.4.

The rollers are 65mm in diameter and were set to rotate at 0.5 revolutions per second so an apple of 67mm diameter would move through one full rotation once every 2 seconds. Although not all apples are exactly 67mm in diameter this is regarded as the industry standard of an average apple diameter and thus the apparatus was designed around this being the average diameter of an apple. At this speed of rotation the camera captures 30 complete images per apple revolution or an image every 12 degrees of rotation. Fruit that is smaller than average would be rotated more than once in the 2s time period which would not cause any issues as the full fruit surface would still be inspected however fruit that was larger than average would not complete a full rotation in 30 camera frames. If required this problem of larger diameters fruit can be solved by tumbling the fruit under the camera for 3 seconds which would allow a very large apple with a diameter of 97.5mm to be fully mapped.

The rollers were powered using a stepper motor driven by a Lam Technologies DS1044 stepper driver. The rollers and stepper motor are connected using a SKF XL series timing belt and pulleys to ensure the rollers rotate at precisely the desired speed.

A single camera placed vertically above the apples path, camera C1 in Figure 3.4, would not be able to capture a complete surface image because of the spherical nature of an apple. For this reason the three-camera arrangement shown in Figure 3.5 was used in order to capture a complete 180 degree view of the apple's surface.

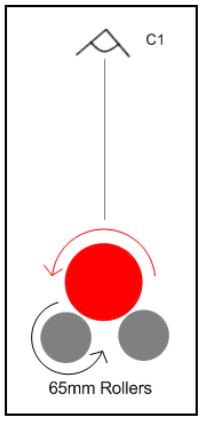

Figure 3.4 Camera and roller arrangement

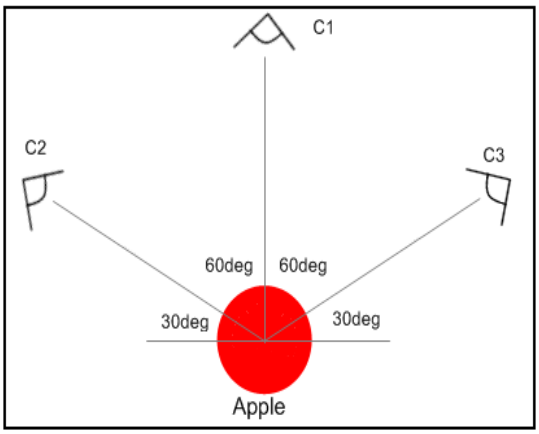

Figure 3.5 Camera arrangement to capture 180 degrees of the apple surface

Cameras C2 and C3 are placed 30 degrees from the horizontal to minimize the separation distance between neighbouring apple lanes (if they were placed any lower, then apples from neighbouring lanes would impede the camera's view of the subject apple), shown in Figure 3.6.

The cameras were mounted at a distance of 340mm from the centre of the apple. This gave a field of view of 450mm x 330mm. An average apple of 67mm diameter had a pixel resolution of approximately 170x170 pixels. Along the longest axis of a pack house conveyor (the direction shown on Figure 3.7), it was estimated approximately 4 apples could be in view at any time and up to two apples on the short axis (orthogonal to direction shown on Figure 3.7). This would allow multiple apples to be processed simultaneously (Figure 3.7) reducing the number of cameras required.

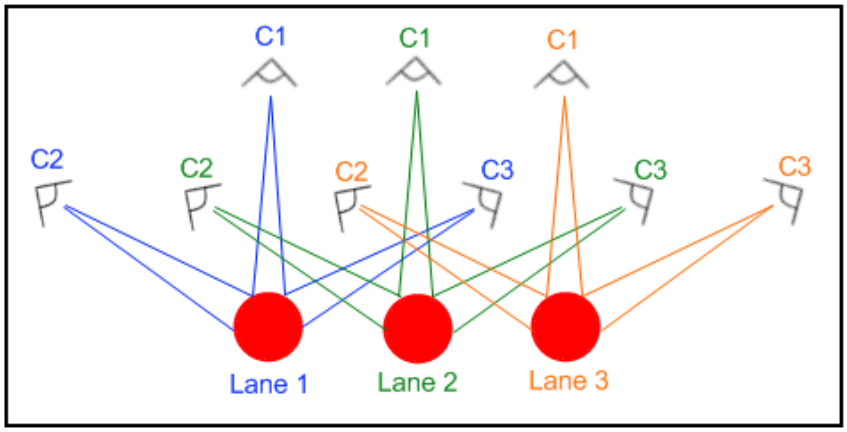

Figure 3.6 Camera arrangement for a pack house conveyor

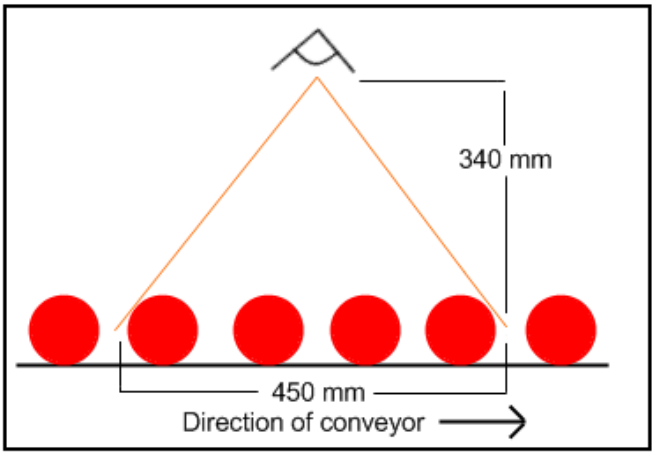

Figure 3.7 Cameras field of view

3.2. Apple Image Library

An image library was created for four varieties of apples; Royal Gala, Braeburn, Eve and Granny Smith. This image library was created to allow detection algorithms to be tested on a known and repetitive set of images allowing for direct comparison between results. Figure 3-8 shows an example of each of these four varieties.

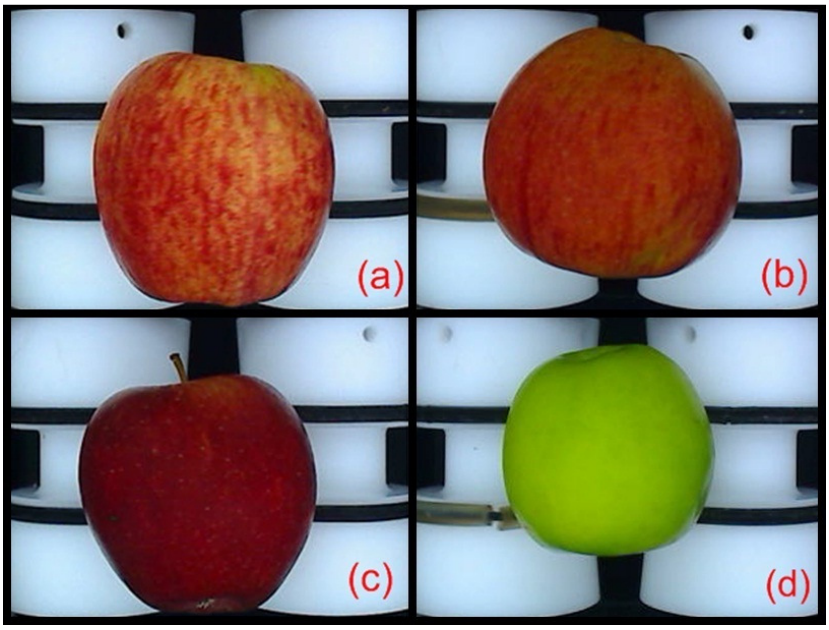

Figure 3.8 Samples of (a) Royal Gala, (b) Braeburn, (c) Eve and (d) Granny Smith

Royal Gala apples are small to medium in size with dark red longitudinal stippling on an orange\green background coloured skin. Braeburn apples are medium to large in size with red stippling on an orange\yellow background coloured skin. The area around the stem tends to be green\brown in colour. Eve apples are medium to large in size with a dark red coloured skin. Granny Smith apples are medium sized with a light green coloured skin and no discernible patterning.

Between 150 and 200 apples of each variety were used to create the apple image library. The apples were a mixture of rejects and good fruit from an industrial packing line. The rejects had a large variety of defects. 30 images were taken for each apple as it was tumbled in two different orientations (stem pointing towards the West and stem pointing towards the North). The inspection algorithms were tested on the resulting 20,880 image library. A summary of the image library is given in Table 3.1. Typical images are shown in Figure 3.9.

Table 3-1 Apples in image library

Apple Variety	Number of Apples	Number of Images
Royal Gala	154	4,620
Braeburn	154	4,620
Eve	198	5,940
Granny Smith	190	5,700
Total	696	20,880

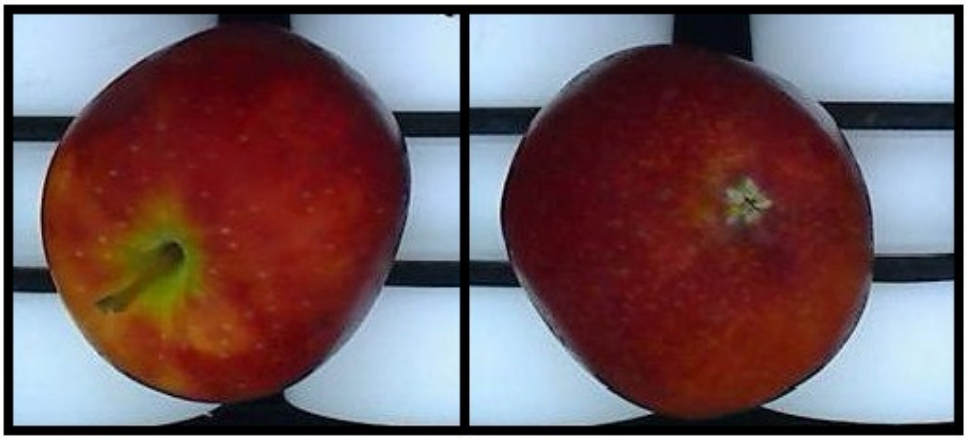

Figure 3.9 Typical images of an Eve apple. Stem is visible on left and calyx is visible on right

The apple library was made up of only colour images. When the webcam was operated in infrared mode, the exposure time of the camera had to be increased to accommodate the low level of light passing through the IR filter. This increased exposure time resulted in blurred images of the rotating apples. Infrared images of stationary apples removed most of the textures from the apples surface and showed imperfection on the skin well but did not hold any additional information over the colour images.

3.3. Random library test sample

A random sample of 2,155 apple images (encompassing all 4 apple varieties) was selected from the apple image library in order to test any inspection algorithms. The sample size was determined using a 95% confidence interval, a 2% confidence level and a library population

of 20,880 images. Having a sample size representing approximately 10% of the total library images allowed for multiple tests to be performed quickly yet still yield representative data. The distribution of stems, calyxes and defects in the random sample, as defined by the stem and calyx table is shown in Table 3-2.

Table 3-2 Random sample population

	Stems	Calyxes	Defects	Total
Total	287	234	230	751

3.4. Stem, Calyx and Defects reference table

Each of the images in the library was inspected manually and any objects of interest within the search area on the apple (for example a stem, calyx, or defect) were noted in a reference table. Also noted in the table were the overall sharpness of the image (occasionally an apple may skip as the rollers rotate which results in a blurred image) and the proximity of an object of interest to the edge of the search area. The results of inspection algorithms could easily be checked against the reference table. The result of a particular image in the library might be negated as not to effect on the overall result for factors such as a blurred image or a defect which was known to be very close to the boundary of the search area.

3.5. Colour spaces

3.5.1. Three-space search (RGB)

The equations allowing the likeness of two points in RGB (red-green-blue) space to be compared are shown below. Equation (2) is used to calculate the angle displacement between the vectors while Equation (3) calculates the difference in length of the vectors. The difference in angle shows how similar the points are with respect to their colour while the

difference in length shows how similar their brightness is. The smaller the difference the more similar the RGB points are.

$$\underline{a}.\underline{b} = |a||b|.\cos\theta \quad (1)$$

Where $\underline{a}$ is the unit vector for the first RGB pixel and $\underline{b}$ is the unit vector for the second RGB pixel. θ is the angle between the two vectors $\underline{a}$ and $\underline{b}$.

$$\Delta\theta = Cosine \frac{(R_1 \times R_2 + G_1 \times G_2 + B_1 \times B_2)}{\sqrt{{R_1}^2 + {G_1}^2 + {B_1}^2} \times \sqrt{{R_2}^2 + {G_2}^2 + {B_2}^2}} \quad (2)$$

$$\Delta Length = \sqrt{R_1^2 + G_1^2 + B_1^2} - \sqrt{R_2^2 + G_2^2 + B_2^2} \quad (3)$$

Where R is the red component, G is the green component and B is the blue component of a pixel in RGB space.

3.5.2. Three-space search (HSL and HSI)

HSL (hue, saturation, and lightness) and HSI (hue, saturation, and intensity) are the two cylindrical-coordinate based representations of points in an RGB colour space.

$$Hue = \frac{60(G-B)}{RGB_{max} - RGB_{min}} \quad if\ RGB_{max} = R \quad (4)$$

$$Hue = \frac{120(B-R)}{RGB_{max} - RGB_{min}} \quad if\ RGB_{max} = G \quad (5)$$

$$Hue = \frac{240(R-G)}{RGB_{max} - RGB_{min}} \quad if\ RGB_{max} = B \quad (6)$$

$$Saturation = \frac{RGB_{max} - RGB_{min}}{RGB_{max}} \quad (7)$$

$$Luminance = 1/2(RGB_{max} - RGB_{min}) \quad (8)$$

$$Intensity = RGB_{max} \quad (9)$$

3.5.3. Two-space search (r', g', b')

The equations to convert from RGB space to 2 space r'g', r'b' or b'g' are shown below.

$$r' = \frac{R}{\sqrt{R^2 + G^2 + B^2}} \quad (10)$$

$$g' = \frac{G}{\sqrt{R^2 + G^2 + B^2}} \quad (11)$$

$$b' = \frac{B}{\sqrt{R^2 + G^2 + B^2}} \quad (12)$$

This allows for an image to be analysed in two-space but still represents the true colour of the pixels. There is very little blue in apples so the points in r'g' space lie on a radius ($\sqrt{r'^2 + g'^2}$) = 1 boundary. The greater the blue component of a pixel, the further the point moves away from the boundary of r'g' = 1.

3.5.4. One-space search (Monochrome)

The equations to calculate the monochrome or greyscale value for a pixel from an RGB image are shown below. The second equation allows for scaling of the three component colours, RGB, so more or less influence can be given to a specific colour.

$$GR = \frac{(R + G + B)}{3} \quad (13)$$

$$GR = \frac{((R \times rf) + (G \times gf) + (B \times bf))}{3} \quad (14)$$

Where rf, gf and bf are factors used to influence the red, green, and blue components on the grey scale value of a pixel. rf + gf + bf = 1.

4. Analysis

In order to find if an apple has any unwanted blemishes, the stem and calyx must first be identified. If the positions of the stem and calyx are known, any other irregularities on the apple skin surface can easily be identified. The size, colour and shape of stems and calyxes vary a great deal between apples and apple varieties. Due to the varying characteristics several techniques have been used to try and identify them. Initially Pixels of Interest (POI) must be found. These POI are any pixels that are not normal apple skin.

4.1. Identifying the apple boundaries

The relative positions of the rollers and the camera are fixed and although the tumbling apple remains approximately in the same place, the centre of the apple must be found in each image so that the inspection algorithm can be centred about this point. The background of the image is black while the rollers are white and the apple is coloured. These contrasting colours allow for the North, East, South and West edges of the apple in the image to be found easily and quickly. Search boundaries are set so that nothing outside the apple is searched which reduces unnecessary computation. Figure 4.1 shows these boundaries, demarcated by small green circles, on a typical image. These boundaries are also used to determine an approximate apple diameter (15).

$$\text{Apple } \emptyset = \frac{(y_{South} - y_{North}) + (x_{East} - x_{West})}{2} \quad \textbf{(15)}$$

4.2. Identifying the area of an apple to be inspected

In an attempt to further minimize the computation time, an apple search area had to be defined. Two shapes of the search area were considered, namely round or square. The inherent problem with a round search area is the computation needed to calculate the XY coordinates of a pixel trigonometrically and the need to ensure that a pixel is not inspected

multiple times. To solve this, a lookup table was created so that, for a known radius, the pixels could be scanned efficiently.

Images from multiple cameras, positioned radially about the plane in which the apple is rotated, are interlaced to create a complete image of the apple's surface. As round search boundaries are interlaced, less of the search areas are overlapped than with square boundaries of equal width. Although some overlap is needed to ensure features do not pass between successive camera's search areas, the overlap represents surface area that is analyzed twice which is a waste of computational power. Thus, from the overlap perspective the round search area was better than the square one. The round search area had one other advantage; the apple shape is approximately spherical which means that light intensity varies radially. In a square search area the variation in light intensity around the perimeter of the square (with its varying radial distance from the centre of the search area) would have a significant and undesirable effect on inspection algorithms. Thus, the round search area was chosen, with a diameter of 50% of the apple diameter. This is shown in Figure 4.1 by the white circle. With the proposed three camera arrangement this search area had a 5 degree overlap between cameras while maximizing the inspected surface per image and minimizing the extra computation required in processing the overlap.

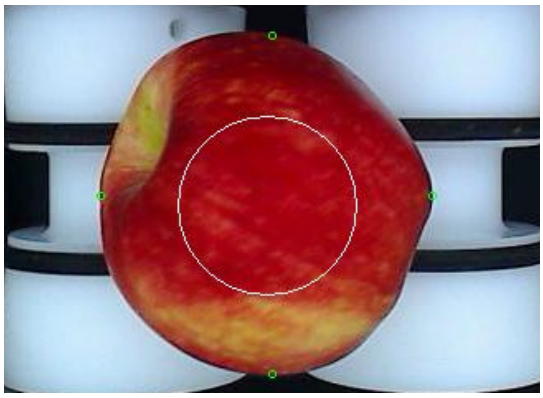

Figure 4.1 Apple image showing search area (inside the white circle) and North, East, South and West boundaries (marked with small green circles)

4.3. Finding pixels of interest using a colour filter

The typical search area of an apple had a radius of 90 pixels and a calyx has a radius of approximately 10 pixels. To reduce the computation a simple colour filter can be applied to distinguish between abnormal and normal pixels on the apple's skin. This filter can greatly reduce the number of points which will need further examination for defects and thus reduces computation times. If, for example, only a calyx and normal skin are within the search area, using typical sizes stated above, the number of pixels being further examined is reduced from approximately 25,500 pixels to 300 pixels or a 98.5% reduction.

Several techniques were developed and are evaluated below to search a specified area for pixel anomalies. These anomalies can be recorded as POI that can later be analyzed. The techniques can be used in any of the colour spaces described in section 3.5 to identify POI visible in that colour space.

Table 4-1 Numbered array representation of pixels in an image

1	2	3	4	5	6
7	8	9	10	11	12
13	14	15	16	17	18
19	20	21	22	23	24
25	26	27	28	29	30
31	32	33	34	35	36

4.3.1. XY Scan

The XY scan runs first horizontally (pixels 1-2, 2-3, etc) and then vertically (pixels 1-7, 7-13, etc) within the search area boundaries looking for a discrete difference between two consecutive pixels, in the desired colour space. This discrete difference or step change between two pixels can be quantified by using a threshold to identify pixels that are different or not.

4.3.2. XY 3 point average scan

The XY 3 point average scan runs horizontally (average pixels[1+2+3]-[4+5+6], etc) and then vertically (average pixels [1+7+13]-[19+25+31], etc), like the XY scan, but instead averages the current pixel with the previous two and compares it to the average of the following three pixels. If there is a large step change between the two averages the current pixel can be identified as a POI.

4.3.3. Averaged search area scan

The averaged search area scan creates a reference point average for the search area by averaging all the pixels in the search area, in the colour space of interest. The scan then compares each point in the search area to the reference-point average and identifies pixels that deviate, more than a predefined threshold, from the reference-point average as anomalies.

4.3.4. Averaged search area 9 point scan

The averaged search area 9 point scan uses the same reference-point average of the search area, as the previous technique but with the average of the pixels based on the 9 points in the immediate vicinity. For example if pixel 8 is the current pixel the average of pixels; 1, 2, 3, 9, 15, 14, 13, 7 and 8 is calculated to give the point average. If the point average is too dissimilar, set by a predefined threshold, from the search area the pixel is recorded as an anomaly.

4.3.5. Averaged search area frame scan

The averaged search area frame scan averages the pixels around the border of the search area to give the search area reference point average. This border average is significantly faster than the previous technique, which averages all the pixels in the search area, because of the greatly reduced number of pixels being averaged. This reference average is more susceptible to being skewed if a large anomaly lies on the border but still gives a good

representation of the average colour. Each pixel in the search area is then compared to this reference average with anomalies being recorded if the difference between the pixels is larger than the allowable threshold.

4.3.6. Averaged search area frame 9 point scan

The averaged search area frame 9 point scans creates a reference border average like the previous technique but instead compares this to a 9 point average of a pixel. This 9 point average is calculated from the pixels in the immediate vicinity of the current pixel, similarly to the averaged search area 9 point scan method. This 9 point average is compared to the frame average and any anomalies identified.

4.3.7. Radius 9 point centre scan

The radius 9 point centre scan uses the same 9 point average about the current pixel but instead compares this 9 point average value to the average of 24 points, on a circle, a set pixel radius from itself. If the difference between centre average and circle average is above a threshold, the pixel is recorded as an anomaly

4.4. RGB, HSL and monochrome analysis

The apple images were analyzed in Red-Green-Blue (RGB) space, Hue Saturation Intensity (HSI) space and Grey Level (Monochrome) space. RGB space proved the most reliable at finding POI that included the stems, calyxes and defects while HSI space found calyx POI well. The difference between the two colour spaces is shown (Figure 4.2 and Figure 4.3) for a Braeburn apple's stem and calyx in RGB and HSI space. The blue fraction of the RGB pixels was ignored as blue is not a colour common in apples and although it can be used to identify calyxes, the calyx was equally visible in the red and green portions of the RGB pixels.

The sample of 100 apple views which had stems, calyxes and defects present were visually inspected to determine which method and threshold identified pixel abnormalities best. Of the

seven methods, sections 4.3.1 to 4.3.7, described for finding abnormal pixels, the averaged search area 9 point scan, section 4.3.4, provided the most reliable results for finding abnormal pixels when tested against a small sample of apples. It was the slowest of the searches because of the large number of points analysed but this could later be optimized to increase the speed. The averaged search area 9 point scan was performed in the red and green portions of the RGB space with pixels that didn't meet the discrimination threshold being identified as abnormal pixels. The two pixel maps of the pixels identified in red and green were compared with any pixels shared between the two colour spaces being recorded as a Pixel of Interest (POI) for further analysis.

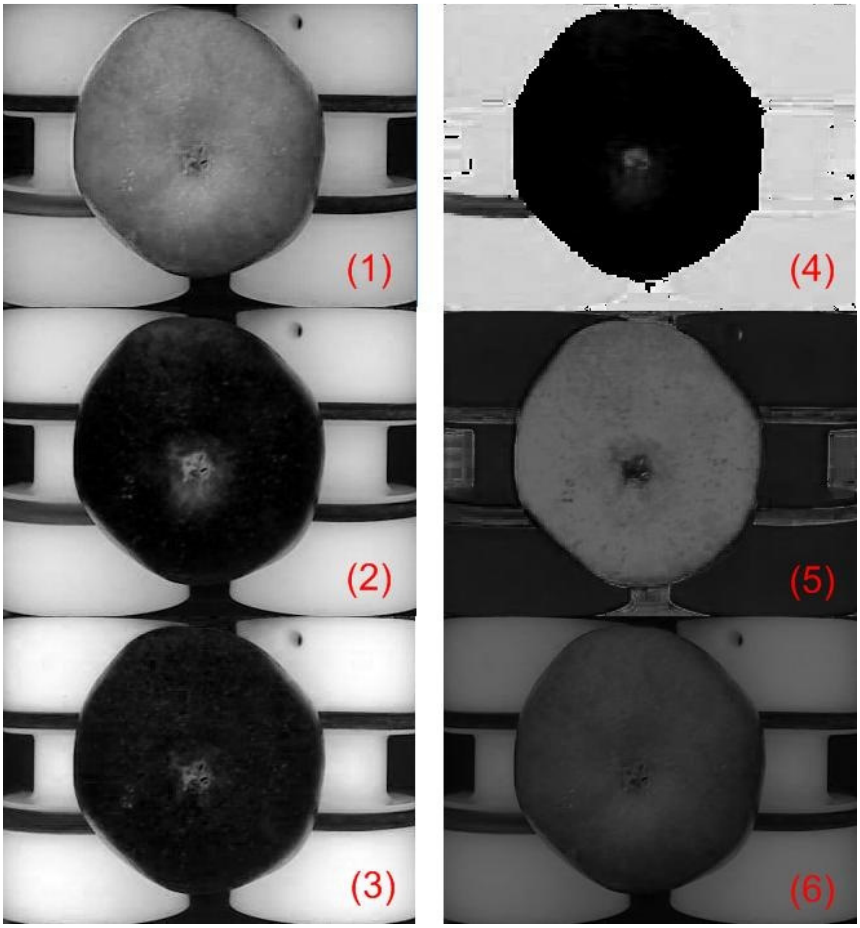

Figure 4.2 Calyx end view of a Braeburn apple showing images for

(1) Red, (2) Green, (3) Blue, (4) Hue, (5) Saturation and (6) Intensity

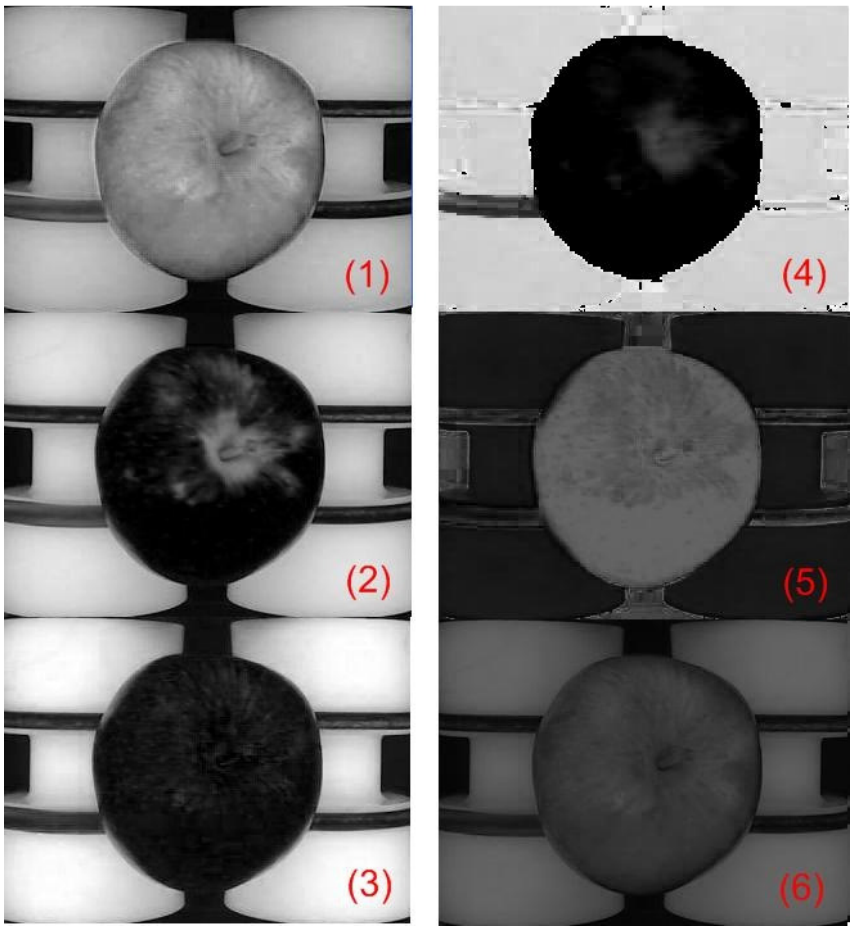

Figure 4.3 Stem end view of a Braeburn apple showing images for

(1) Red, (2) Green, (3) Blue, (4) Hue, (5) Saturation and (6) Intensity

4.4.1. Sobel filter

A Sobel type filter was applied to any POI-flagged pixels found by the previous colour filter. The filter takes the POI and compares its greyscale value to the greyscale value of pixels in its vicinity at a set radius, shown in Figure 4.4. The radius of the Sobel comparison was set at 5 pixels, which was approximately half the size of an average calyx. If the change in greyscale is above a threshold, and the change is from light to dark, the POI remains a POI, else it is removed from the list of POI. The change must be from light to dark because the stem/calyx tends to have a lighter greyscale value than the surrounding skin.

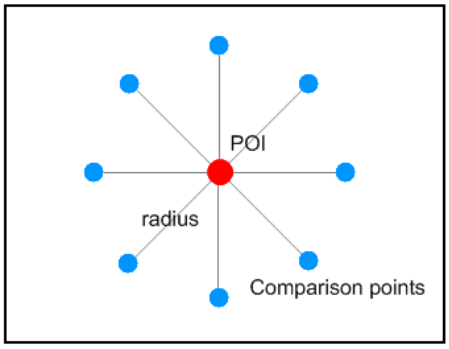

Figure 4.4 Points used for the Sobel filter

4.4.2. Islanding filter

The POI identified by the colour and Sobel filter are gathered into islands using an 'islanding' routine also known as a connected component labelling routine. This routine discards isolated pixels that are not part of the main island and ensures the island has a population large enough to be an island of interest.

The POI were gathered into islands by the islanding routine. At this point 8.7% (65/751) of the total sample where there were islands present, as defined by the stem, calyx and defect reference table (Section 3.3), remained unidentified. This was 3.1% (65/2155) of the total views in the random sample. 17% of the defect islands were missed because some apple defects could not be reliably detected by the initial colour filter. An island was incorrectly identified in 72/2155 views where the stem, calyx and defect reference table did not identify one as being present. These were labelled as false positives.

Table 4-2 Random sample population

	Stems	Calyxes	Defects	Total
Identified Islands	266	229	191	686
Unidentified Islands	21	5	39	65
Totals	287	234	230	751

4.4.3. Island characteristics filter

These islands have several useful characteristics, such as population, density, aspect ratio, radius of gyration etc., that can be used to determine if a stem or calyx is present. These characteristics and their significance are discussed below.

1. The population, Pop, is simply the number of points in the island. Upper and lower population threshold limits are set to ensure an island has a similar number of points to a stem or calyx. A sample of 50 apples per apple variety was used to determine the average size of a stem or calyx island. The population of the stem and calyx for each of the apple's was measured and compared as a function of the apples radius. The upper and lower threshold limits are given in Equation (16) and Equation (17) where r_{apple} is the radius of the apple in pixels.

$$Pop_{lower\ limit} = \pi \left(\frac{r_{apple}}{20}\right)^2 \quad \textbf{(16)}$$

$$Pop_{upper\ limit} = \pi \left(\frac{r_{apple}}{6.5}\right)^2 \quad \textbf{(17)}$$

2. The Centre of Gravity (COG) of the island in the x- direction and y-direction, COGx and COGy are determined using Equation (18) and Equation (19).

$$COG_x = \frac{\sum_{n=1}^{Pop}(x_i)}{Pop} \quad \textbf{(18)}$$

$$COG_y = \frac{\sum_{n=1}^{Pop}(y_i)}{Pop} \quad \textbf{(19)}$$

3. The first moment of area, $J_{first}xy$, is calculated using Equation (20).

$$J_{first}xy = \sum_{n=1}^{Pop} \sqrt{(x_i - COG_x)^2 + (y_i - COG_y)^2} \quad \textbf{(20)}$$

4. The second moment of area, $J_{second}xy$, is a property of the island shape that can be used to determine its resistance to torsion. Similar to $J_{first}xy$, the further the points are from the COG, the higher $J_{second}xy$ becomes. $J_{second}xy$ is calculated using Equation (21).

$$J_{second}xy = \sum_{n=1}^{Pop} (x_i - COG_x)^2 + (y_i - COG_y)^2 \quad \textbf{(21)}$$

5. The Radius of Gyration (ROG) is a property of the island shape that describes the distribution of the area of the island about its COG. The ROG is calculated using Equation (22).

$$ROG = \sqrt{\frac{J_{Second}xy}{Pop}} \quad \textbf{(22)}$$

6. The Radius of Circle (ROC) with equivalent area is the size of the equivalent island if it was a perfectly round disk with an area equal to the current population size. ROC is calculated using Equation (23).

$$ROC = \sqrt{\frac{Pop}{\pi}} \quad \textbf{(23)}$$

7. Aspect ratio (AR) is the ratio of the length between the furthest two points in the island, AR_h, and the width between the furthest two points perpendicular to this length, $AR_{w1} + AR_{w2}$ (Figure 4.5).

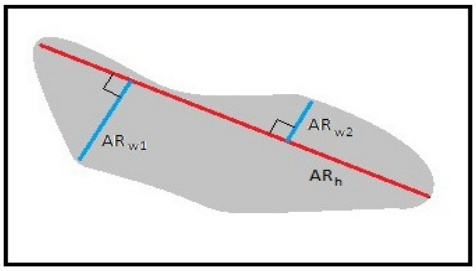

Figure 4.5 Aspect Ratio calculation

8. The density of the island is determined from AR_h and AR_w and is an indication of how solid the island is. It is calculated using Equation (24).

$$\rho = \frac{Pop}{AR_h \times AR_w} \tag{24}$$

The effect of shape on each of the characteristics can be demonstrated using a disc and line constructed from approximately equal pixels (325 pixels for the line and 315 pixels for the disc) shown in Figure 4.6. Table 4-13 summarises the computed characteristics for the two shapes. There is a clear difference between the line and the disc when comparing the first moment of area, second moment of area, radius of gyration and aspect ratio. There is very little difference in the equivalent circle radius and the density because the populations are equal and the shapes are solid.

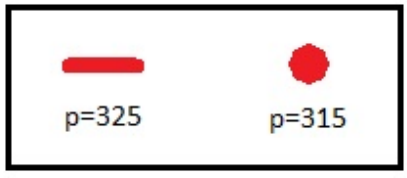

Figure 4.6 Shapes used in island specification tests

Table 4-3 Island characteristics test results for shapes shown in Figure 4.6

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Line	325	3,419	50,142	12.421	10.165	1:0.818	1:0.240	0.792
Disc	315	2,084	17,922	7.543	10.007	1:1.327	1:1	0.787

The red area on a Granny Smith apple shown in Figure 4.7, are the pixels of interest identified by the colour and Sobel filters. The island filter was performed on these POI with the resultant characteristics listed in Table 4-4. The stem and calyx have similar populations but there is a very clear difference between the two when comparing the first moment of area, second moment of area and the aspect ratios. The island characteristics found in Figure 4.7

are not exactly the same as the disc and line characteristics in Figure 4.6 because the stem is shorter and wider than the line and the calyx is more oval than the disc. The disc and line example shows the best case difference in stem and calyx characteristics.

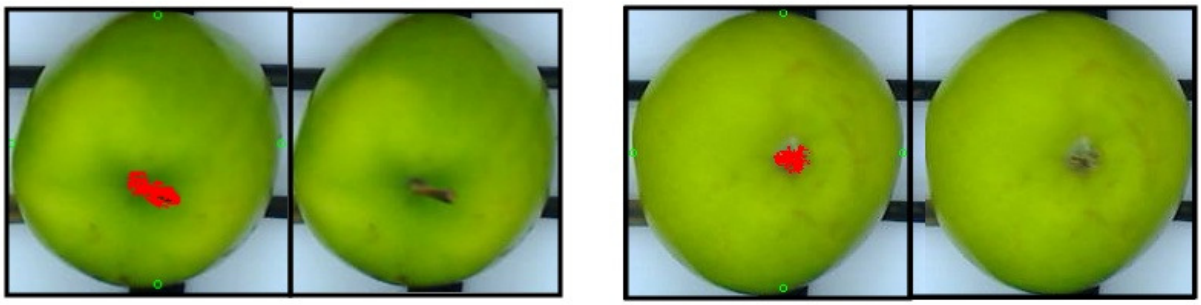

Figure 4.7 Test of characteristics on a Granny Smith's stem and calyx

Table 4-4 Island characteristics test results for Granny Smith apple shown in Figure 4.7

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Stem	336	2,961	35,151	10.228	10.335	1:0.287	1:0.410	0.706
Calyx	285	1,893	16,815	7.68	9.519	1:0.270	1:0.790	0.625

An apple stem can also have been pulled out during the picking process. Although this is undesirable, because it can tear the apple flesh which causes the apple to rot very quickly, it is not a defect. A pulled stem appears as a small circular black dot with a typical radius of approximately 10 pixels (Figure 4.8). The POI found create a halo around the pulled stem because the Sobel part of the algorithm looks for the grey level step change that does not occur around the centre of the dot (which is uniformly grey). The island characteristics for the pulled stem are shown in Table 4-5. Although a pulled stem and a calyx may exhibit several very similar specifications, because of their similar shapes, there is a very large difference in populations which make the two easily distinguishable. There is also a large difference in the shapes of a pulled stem and a non-pulled stem so that it is easy to distinguish between them.

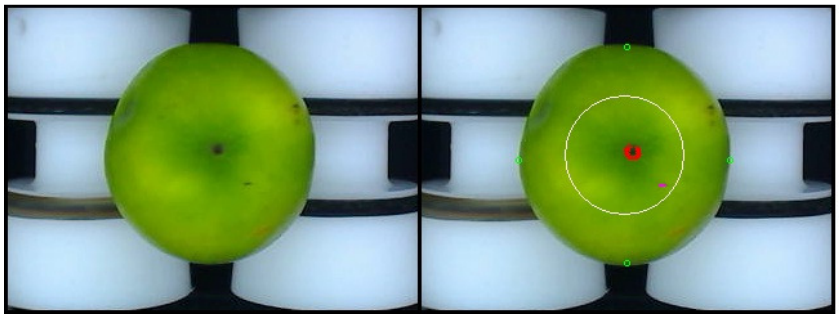

Figure 4.8 Test of characteristics on a Granny Smith apple with a pulled stem

Table 4-5 Island characteristics test results for Granny Smith apple shown in Figure 4.8

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Pulled Stem	87	421	2,625	5.494	5.259	1:0.418	1:0.7	0.621

Leaves that have remained attached to an apple's stem during the picking process are considered a defect because they dry out and do not look aesthetically pleasing when on display in shops. If the apples are being packed by robots leaves can interfere with the robot's ability to pick up the fruit so these apples must be rejected before the packing line.

The colour and Sobel filter finds a large number of points when a leaf is present (Figure 4.9). The characteristics for a leaf are not calculated because the island population exceeds the upper population limit (Equation 2) so the fruit would be rejected from the sorting line. The islanding routine takes too long to calculate the characteristics because of the very large number of POI.

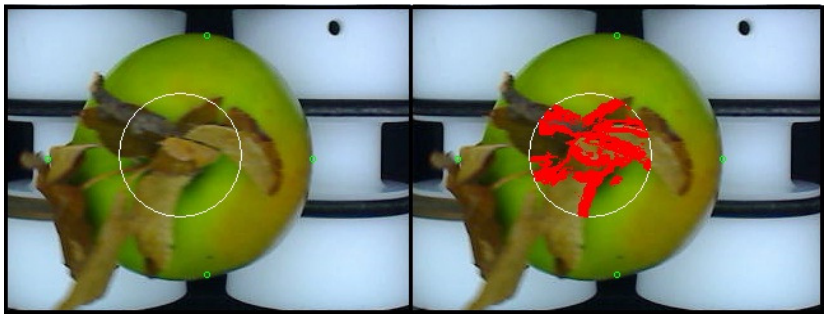

Figure 4.9 Granny Smith apple with leaves present

Several typical apple defects are described below together with their colour and Sobel filter results. These defects can be easily confused with a stem, pulled stem or calyx when comparing their island characteristics.

A stem puncture is a small skin breakage or tear caused by an apple stem being pressed into itself or another apple. These injuries typically occur in the orchard when the apples are placed roughly in the large transport bins by the fruit pickers or when the apples bounce around inside the bin during transport to the packing shed. The injury starts to rot very quickly because of the broken skin. These injuries can easily be misidentified as a stem because the imprinted shape is that of a stem. Figure 4.10 shows a pulled stem, identified in red, and a stem puncture injury, identified in purple. The characteristics of the stem puncture injury are very similar to that of a normal stem, as shown in Table 4-6 and Table 4-4 respectively.

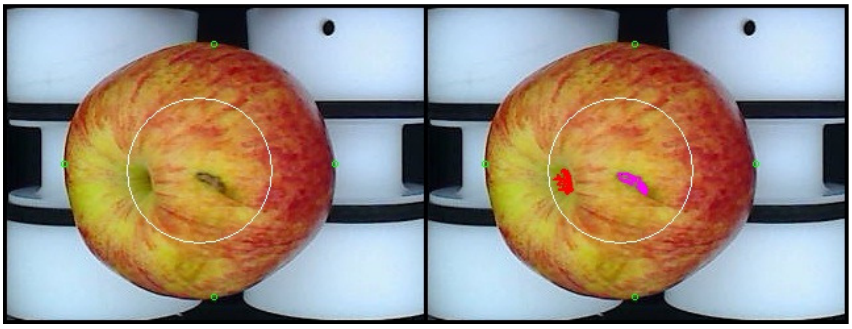

Figure 4.10 Royal Gala apple with a pulled stem (marked red) and a stem puncture injury (marked purple)

Table 4-6 Island characteristics test results for Royal Gala apple shown in Figure 4.10

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Stem Puncture	146	1,009	9626	8.119	6.813	1:0.385	1:0.410	0.691

Bruises are classified as either a new bruise or an old bruise. Old bruises tend to occur during the thinning process which happens 3 to 4 months before the apples are harvested. Thinning is performed to make room for apples to develop fully and to expose them to sunlight.

Sunlight ensures that the apple will develop the correct colour. If the tree is not carefully thinned then as the pruned fruit is cut loose, it can fall down through the tree, hitting and bruising other apples. An old bruise does not show any discolouration but instead an indentation forms were the bruise occurred. A new bruise tends to occur during the picking process when the pickers dump their apples roughly into the bins. These bruises are only hours or days old when the apples reach the packing shed. They are brownish in colour and the flesh under the bruised portion is very soft. The colour and the Sobel filter did not identify pixels of interest when a bruise was present in the image because of the very slight change in colour deviation. This meant that no characteristics could be calculated without adjusting the filter thresholds.

Russet is a yeast-like fungus that affects the skin of apples. It creates a lattice of light brown discoloration. While a small patch of russet can look like a stem or calyx, it typically occurs in large patches. The Russet island identified below in red (Figure 4.11), exhibits island characteristics, shown in Table 4-7, similar to that of a large stem. The island has only just passed the upper population limit criterion for this apple.

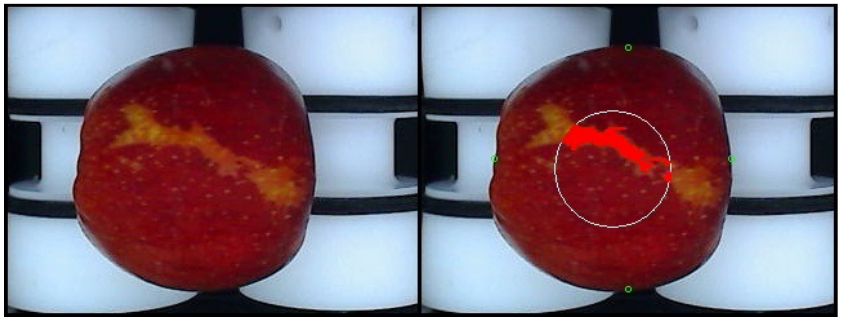

Figure 4.11 Eve apple with a Russet skin discolouration

Table 4-7 Island characteristics test results for Eve apple shown in Figure 4.11

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Russet	915	17,710	467,179	22.596	17.056	1:0.258	1:0.350	0.354

Although apples are washed before being sorted for packing, they may still have patches of dirt. A single small patch of dirt can be misidentified as a pulled stem when comparing island characteristics (Table 4-8 and Figure 4.2). Dirty apples typically have several small patches of dirt and the identification of multiple individual islands distinguishes them from a pulled stem Figure 4.12 shows that the colour and Sobel filters have identified two patches of dirt although only the island characteristics of the larger island are calculated because the smaller island does not meet the island lower population criterion.

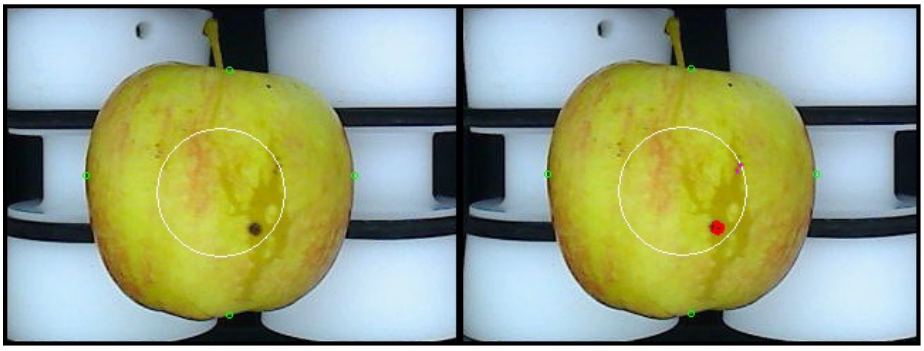

Figure 4.12 Royal Gala apple with dirt spots

Table 4-8 Island characteristics test results for Royal Gala apple shown in Figure 4.12

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Dirt	69	238	1,204	4.178	4.68	1:0.396	1:0.900	0.766

Sunburn or sunscald is caused either by exposure to solar radiation or exposure to freezing conditions, both of which cause damage to the apple skin (Figure 4.13). The affected area of an apple tends to be round in shape and have a light pink colouration. The characteristics of the affected area are similar to that of a calyx although sunburn or sunscald areas are typically much larger than an apple calyx so the island typically does not pass the population upper limit criterion.

Figure 4.13 Braeburn apple with sunburn/sunscald damage

4.4.4. Island characteristics filter

The island characteristics for the random sample were computed so a set of boundaries could be identified to discriminate between the different stem, calyx and defect island classifications. The mean and standard deviation for each island characteristic were calculated from the dataset.

In order to distinguish between stems, calyxes and defects, threshold boundary conditions were set at +/- 2 standard deviations from the mean for the island data sets. This boundary condition meant that statistically approximately 95% of the islands would fit their appropriate classification and are listed in Table 4-9. A voting system was used to determine which category; stem, calyx or defect, an island belonged to. An island got a vote towards a category each time the island characteristic fitted the set of defined boundaries for a category; stem, calyx or defect. The island was identified as the category with the most votes or, if it had not met a minimum number of votes, initially set at 4/8, the island was labelled a defect. The results from the islanding routine were then compared with the stem/calyx/defect reference table, to determine whether an identified island was correctly characterised.

Table 4-9 Island characteristics boundaries, +/- 2 standard deviations from the mean

	Population	$J_{first}xy$	$J_{second}xy$	ROG	ROC	ROG:ROC	AR	Density
Stem	x<1227	x<20136	15,006<x<60,601	6.54<x<13.16	4.70<x<20.92	0.15x<0.42	0.21<x<1	0.16<x<0.72
Calyx	x<948	x<13332	8,281<x<17,506	4.84<x<7.07	2.20<x<18.46	0.09<x<0.57	0.21<x<1	0.07<x<0.71
Defect	x<841	x<12173	1,936<x<59,536	2.35<x<13.05	x<17.24	0.11<x<0.68	0.10<x<1	0.03<x<0.80

It is clear from the table of island characteristics boundaries, Table 4-9, that there is a large overlap between the boundaries of the three different island classifications which will make classification difficult. Table 4-10 shows the results from the classification test using the characteristic boundaries in Table 4-9. True Positive means that the algorithm has identified the island correctly while False Negative means that the algorithm has not identified the island correctly (for example, in the 'Stem' category, there were a total of 266 stems of which 30 were not identified). The category called 'Multi' was used when there was a tie in the voting system.

Table 4-10 Island classification boundaries test results

	True Positive	False Negative	Total
Stem	236	30	266
Calyx	191	38	229
Defect	180	11	191
Multi	539		

The number of false negatives recorded for the categories stem, calyx and defect ranged from approximately 5.8% to 16.6% of the total number for that category. The number of islands that could have been misclassified is represented by the "Multi" value, Table 4-10, which comprises of the number of times there was a tie in leading vote. When this occurred the vote went to the one that corresponded to the stem/calyx/defect reference table but approximately 78.6% of total sample population tests resulted in a tie between more than one of the classifications. This shows the overlap between the island characteristic boundaries is too high to permit accurate classification.

A regression algorithm was used to find the boundaries giving the maximum number of True Positives and minimum False Negatives. This regression algorithm allows for the relationship between the dependant variable and the independent variables to be analysed by varying the

value of one of the independent variables at a time and noting its effects on the dependant variable. In this case the number of True Positives was the dependant variable and the eight boundary thresholds were the independent variables. This allowed for the optimum boundary relationships to be quickly identified.

The optimum boundary was +/- 2.5 standard deviations and the results using these boundaries are given in Table 4-11. The upper limit of the regression threshold was set as 2.5 standard deviations and because the regression was not being punished for a high number of multi votes, the algorithm moved towards the thresholds that would allow most of the islands to fit. The total number of False Negatives decreased to 2.6% but the tied votes are unacceptably high (95.8%).

Table 4-11 Island classification boundaries test results from regression algorithm

	True Positive	False Negative	Total
Stem	261	5	266
Calyx	219	10	229
Defect	188	3	191
Multi	657		

4.4.5. Concentric ring filter for colour graduations

Due to the poor results from the direct analysis of the island identified by the colour and Sobel filters, a new method that focused on the area surrounding the island was devised. The method analyzed the change in light intensity moving in concentric rings radially from the island COG (Figure 4.14). The colour space of each ring is averaged and compared to the other rings; r1, r2, r3.

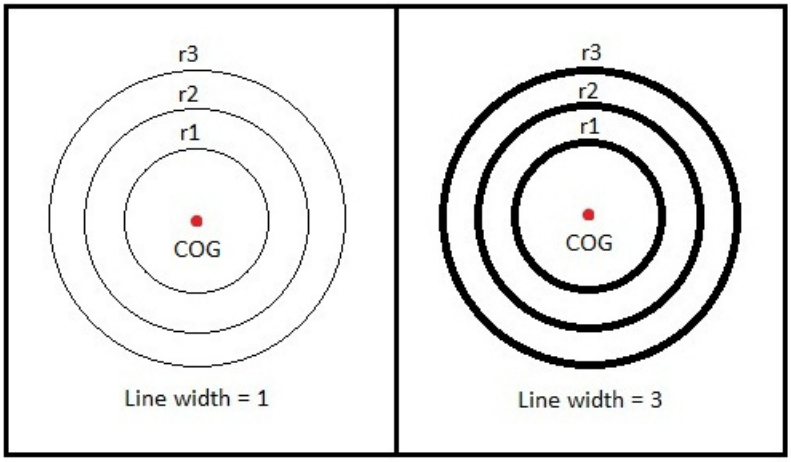

Figure 4.14 Concentric ring filter, single pixel width average vs. 3 pixel width average

The change in curvature of an apple surface is greater in the stem and calyx regions. This implies that the difference in light intensity between concentric rings around a stem or calyx island will be greater than between concentric rings around a defect island.

Different colour spaces (greyscale, RGB and HSI) were used to analyze the concentric rings. The pixels at a set radius, (r1, r2, etc.) are averaged to form a ring average. This ring average is then plotted against its distance radially from the island's COG. The separation between rings was manipulated to maximize this effect along with averaging several consecutive radii to capture a wider band of data to form a single ring, Figure 4.14. A wider band average gave a slightly less noisy plot while changing the separation between the rings just increased the resolution of the graph. Changing neither made significant improvements to the graph so a single pixel ring was used to average 10 concentric rings with a ring separation of 5% of the apple radius so as to maximize the speed of the routine. The island COG_{xy} is shown as a red dot while the concentric rings are plotted in white (Figure 4.15 and Figure 4.16).

During initial testing it was determined that the Red component of the RGB space was best suited for red apples such as Royal Gala (Figure 4.15) while the Green component of the RGB Space was best suited for green apples such as Granny Smiths (Figure 4.16). The

Intensity component of the HSI space gave the most relevant results overall when a dimple was present because there is direct relationship between the prominent colour of an apple and the intensity values. The calculated intensity value of a pixel is equal to the maximum of the three RGB components so when an apple is predominantly red, the Intensity plot is directly related to the RGB's Red component while with a predominantly green apple, the Intensity plot is directly related to the RGB's Green component (Figure 4.15 and Figure 4.16).

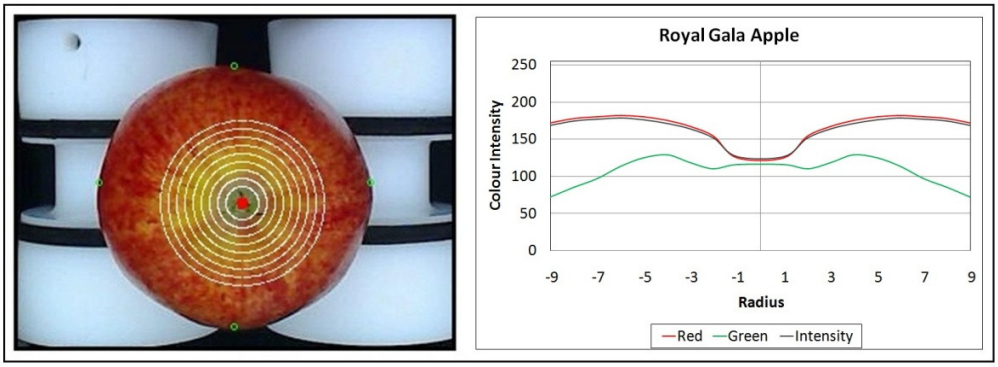

Figure 4.15 Concentric ring plot for Royal Gala apple calyx

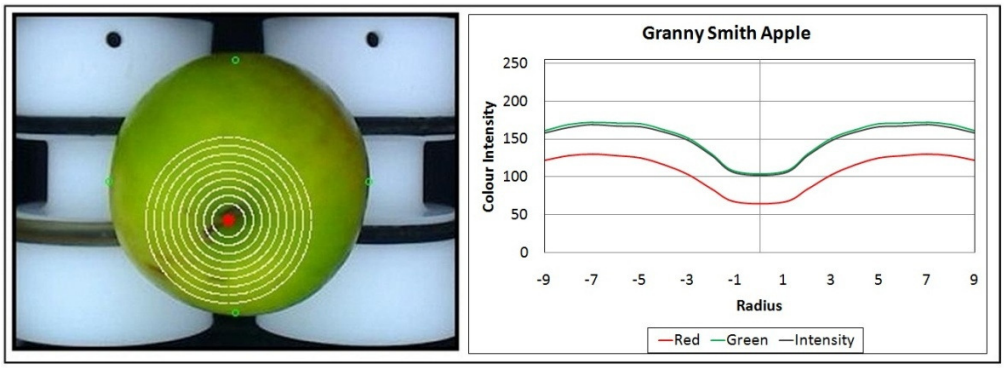

Figure 4.16 Concentric ring plot for Granny Smith apple stem

The intensity profiles for a stem, calyx and defect were plotted and overlaid on the same graph for a red and a green apple, Royal Gala apple (Figure 4.17) and Granny Smith apple (Figure 4.18). It is evident that the stem and calyx profiles are similar while the defects profile is significantly different from the stem and calyx. The three profiles confirm the concept of the light intensity following the shape of the apple. The stem and calyx profiles highlight the

concave shape of the apple that occurs near a stem or calyx while the defect profile demonstrates a smooth convex shape except for the dimple where the defect is.

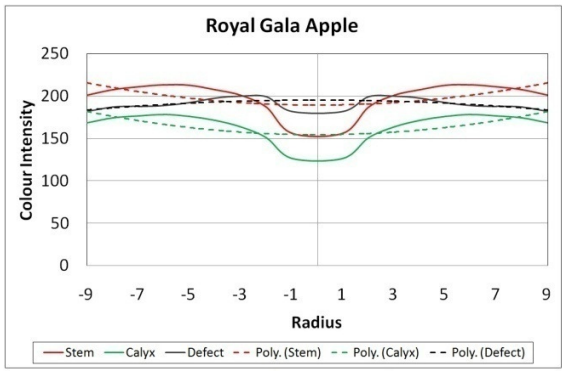

Figure 4.17 Royal Gala intensity profiles

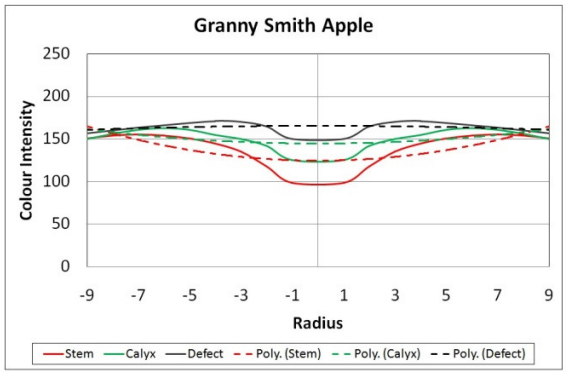

Figure 4.18 Granny Smith intensity profiles

An algorithm was designed to distinguish between islands that appear to have occurred at the concave ends of the apple as opposed to islands on the normal convex area. A second order polynomial line, illustrated by the dotted line in Figure 4.17 and Figure 4.18, was fitted to the data set which could distinguish between the two conditions. If the fitted polynomial is positively signed the island dataset is in a concave area of the apple while if the polynomial is negatively signed the data set represents an island on a convex area. The radius at which the maximum intensity occurs from the centre of the island can also indicate whether the area is concave or convex in shape. A second algorithm finds this point where the profile fitted to the data set gradient changes from being positive to being negative as it moves outward through the concentric rings. This gradient change defines the radius where the maximum intensity of the dataset occurs. If the island is in a concave area the maximum intensity occurs at a greater radius from the island COG than when the island occurs on a convex area. The maximum intensity of the defect data set illustrated in Figure 4.17 occurs at a radius of 2 while for the same apple, the stem and calyx data sets maximum occurs at a radius of 6.

An initial test using the two filter algorithms for fitted line gradient and maximum point radius was run against the library random test sample. A fixed threshold was used to decide whether

an island occurred in a concave area or on a convex area of the apple. If the fitted line was positive and had a maximum point radius greater than 3 the island was classified as a stem/calyx while if the fitted line was negative and the maximum point radius was 3 or less the island was classified as a defect. These classifications were then compared to the stem, calyx and defect reference table for the random library sample set, Table 4-12.

Table 4-12 Test results of random library sample for initial thresholds

	Stem/Calyx	Defect
True Positive	363	51
False Negative	132	140
Total	495	191

To improve these results from the initial test a regression algorithm was run to optimize the thresholds so as to reduce the number of false negatives. The regression algorithm negated the use of the maximum point radius because it caused a large number of the false negatives and thus skewed the results. In the regression algorithm an equal weighting was given towards the algorithm getting a stem, calyx or defect correct which resulted in 81% of the islands being correctly identified (Table 4-13). This is a significantly better result than the island characteristics testing because there is a clear distinction between the two classifications, stem/calyx or defect, however this level of accuracy would still not be acceptable for a commercial application.

Table 4-13 Test results of random library sample for optimized thresholds

	Stem/Calyx	Defect
True Positive	400	155
False Negative	96	36
Total	495	191

4.4.6. Finding islands using concentric ring filter

The previous two methods, Island characteristics and Colour graduations, for identifying stems, calyxes and defects relied on the initial colour and Sobel search to find islands of interest off which the concentric ring filter could be run. This colour filter only found 91.3% of the total stems, calyxes and defects present in the sample population so before the concentric ring filter started it already had an error of 8.7% of the sample population. In an attempt to improve the accuracy with which islands are found, the centre of the concentric rings was set at the centre of the apple. The intensity curve of the averaged rings was plotted against radius and both error and sum of the absolute error squared between the fitted second order polynomial and the plotted points were examined. The sum of absolute errors squared is included because it should allow smaller islands to be detected that would otherwise only create a small error.

Figure 4.19 shows the intensity profile and fitted polynomial for a Granny Smith apple with defect. The small deviation of the intensity profile from the polynomial has a small error, but a large error^2. Figure 4.20 shows the intensity profile and fitted polynomial for a Royal Gala apple where there is no imperfection (stem, calyx or defect). Figure 4.21 shows the intensity profile and fitted polynomial for a Royal Gala apple's calyx. Both the error and the squared error are large. Table 4-14 shows a comparison of the errors for the three different cases.

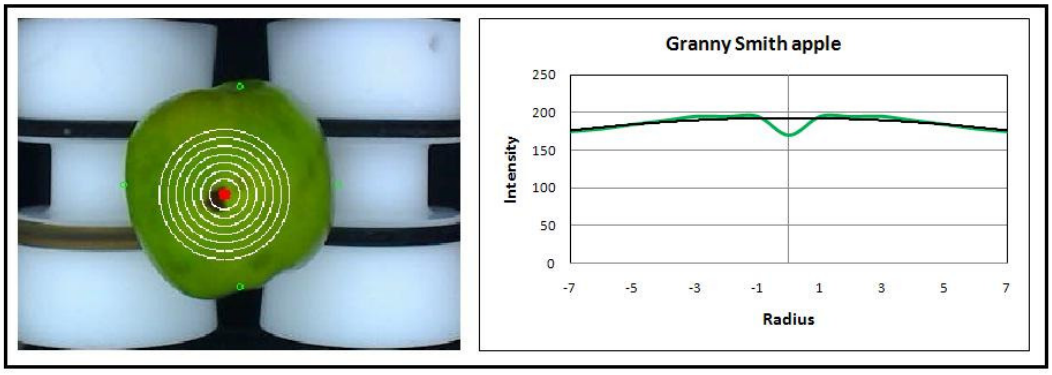

Figure 4.19 Intensity profile (green) and fitted polynomial (black) of Granny Smith apple with defect

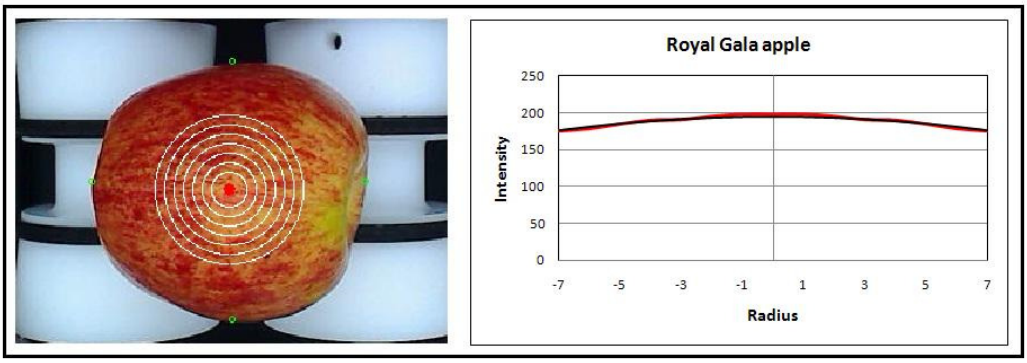

Figure 4.20 Intensity profile (red) and fitted polynomial (black) of Royal Gala apple with no imperfections

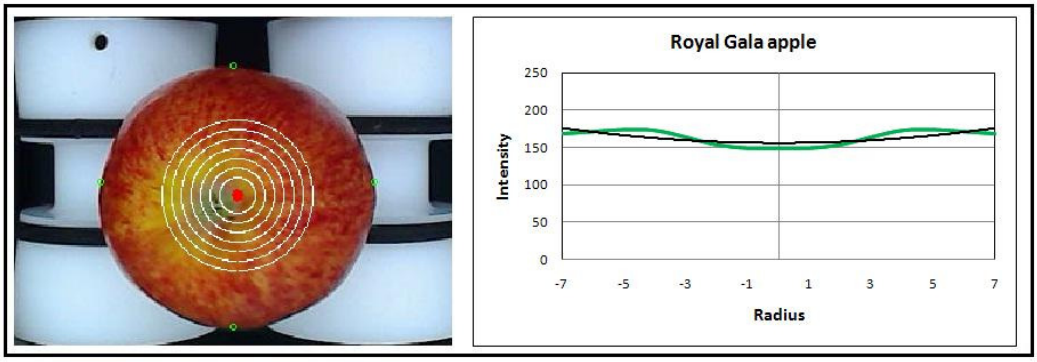

Figure 4.21 Intensity profile (green) and fitted polynomial (black) of Royal Gala apple with object of interest in view

Table 4-14 Error chart for intensity profiles

	error	error^2
Figure 4.19	46.21	641.40
Figure 4.20	12.63	17.73
Figure 4.21	79.73	570.53

This error filter was run against the random library sample and compared the stem, calyx and defect table, with the regression algorithm determining the best thresholds to set error and error^2. For an island of interest to be present thresholds of (error > 40) and (error^2 >100) gave the best results, Table 4-15. The error gave the best result with only 75.6% of the islands correctly identified as being present, which was significantly worse than the results from colour and Sobel filter.

Table 4-15 Test results from error calculations between dataset and fitted polynomial

Stem/Calyx/Defect	error	error^2
True Positive	568	525
False Negative	183	226
Total	751	751

Fitting a second order polynomial to the data set did not identify the islands well because the fitted polynomial could fit the data set well (thus having a low error) but indicated there was a possible island present because of the shape of the fitted polynomial. As with the previous method a positive fitted polynomial indicated a concave area on the apple was present which this method did not account for.

To overcome this problem a "fixed normal skin curve" was created that represented apple skin on a convex area of the apple with no defects present illustrated by the black line Figure 4.22 and Figure 4.23. Equation (26) was formulated from a sample of 50 apples using the average gradients, component A in Equation (25), of the equations for the fitted second order polynomials of the sample population. The normal skin curve has no x offset, component B=0 in Equation (25), because the plots are symmetrical about radius = 0. The y offset remains a variable because the intensity level of each intensity profile changes and this allows the profile to be adjusted to fit the current intensity profile. The y offset is set so the end points, radius = 7, of the intensity profile and the fixed normal profile coincide at radius = 7.

$$\text{Second order polynomial} = A(x)^2 + \mathrm{B}(x) + \mathrm{C} \quad \textbf{(25)}$$

$$\text{Fixed normal skin curve} = 0.4(r)^2 + \mathrm{y} \quad \textbf{(26)}$$

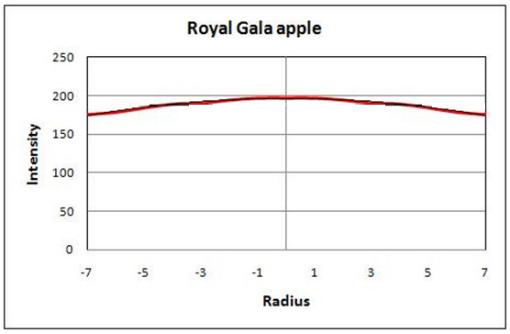

Figure 4.22 Fixed skin curve (black) fitted to normal skin data (red)

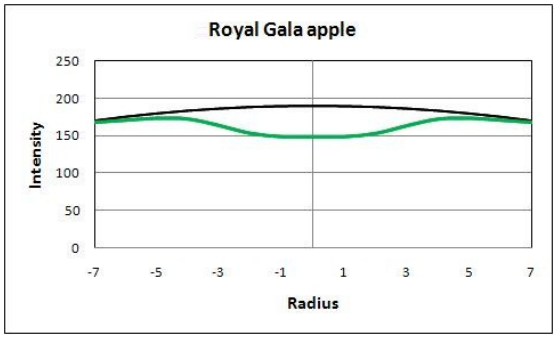

Figure 4.23 Fixed skin curve (black) fitted to calyx area (green)

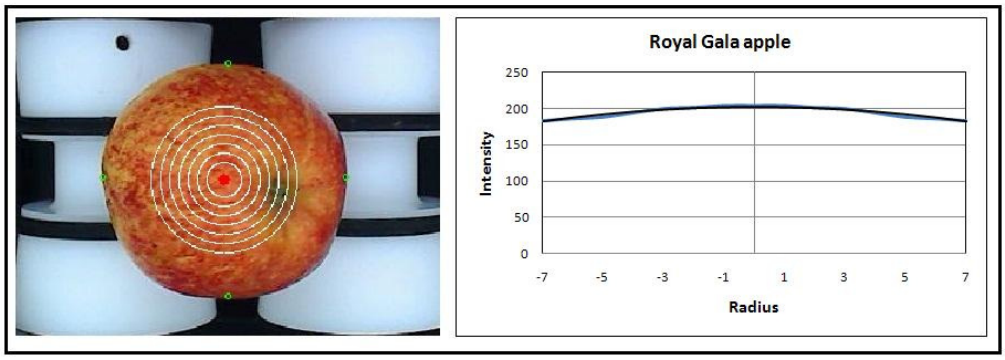

Figure 4.24 Intensity profile (black) of Royal Gala apple failure

The error between the fixed normal curve, Equation (26), and the plotted data set was calculated for two intensity profiles shown in Figure 4.22 and Figure 4.23. The resulting errors, Table 4-16, show a clear distinction between the two cases, when an object of interest is present (Figure 4.23) and when one is not (Figure 4.22). The sum of error squared shows the biggest difference with Figure 4.23 error^2 being 240 times larger than Figure 4.22. The regression algorithm was used to determine the best thresholds for the error cut-off between normal skin and an object of interest being in view.

Table 4-16 Error chart for intensity profiles

	error	error^2
Figure 4.22	19.55	36.26
Figure 4.23	283.57	8657.96
Figure 4.24	20.62	48.74

The regression algorithm identified the best thresholds as (error < 327) and ($error^2$ < 6370) being normal skin otherwise there is an object of interest present. The results using these thresholds for error and $error^2$ for the random library sample were poor, Table 4-17. The $error^2$ threshold gave the best results but only correctly identified 68.18% of the stems, calyxes and defects.

Table 4-17 Test results from error calculations between dataset and fixed curve

Stem/Calyx/Defect	error	$error^2$
True Positive	487	512
False Negative	264	239
Total	751	751

This method where the centre of the concentric rings is placed in the centre of the apple has failed because as the island of interest moves towards the edge of the search boundary its effect on the average of a particular ring's intensity is reduced because of the increased number of pixels in the ring. Figure 4.24 shows a calyx near the edge of the search boundary with its intensity profile adjacent. In the intensity plot a small deviation occurs at radius = 5, which corresponds to the position of the calyx but its effect on the ring's average is minimal. The error and $error^2$, Table 4-16, between the intensity plot and the fixed normal curve for Figure 4.24 fall within the threshold limits, set by the regression algorithm, for no island of interest to be present.

4.4.7. Multiple origin concentric ring filter

A new method using the previous concentric ring filter (Section 4.4.6) with the fixed error comparison filter was created to try and improve the accuracy of the filter by reducing the problem caused by the islands of interest being near the edge of the search boundary. The new method used a series of fixed origins spread symmetrically around the centre of the apple (point 1 in Figure 4.25). A total of 17 origins, represented by red dots in Figure 4.26 and numbered in Figure 4.25, were used including the centre of the apple. The radii r1 and r2,

Figure 4.25, represent the distance from the centre of the apple to the origins. The outer search radius, r2, and the outer radius of the concentric rings filter were set at 1/3 the apple radius. This meant the maximum distance a point could be from the centre of the apple was 2/3 of the apple radius which ensured none of the background image was processed by the filter.

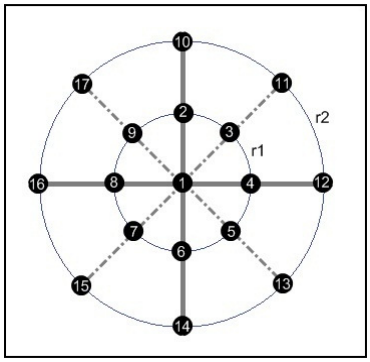

Figure 4.25 Multiple origin concept

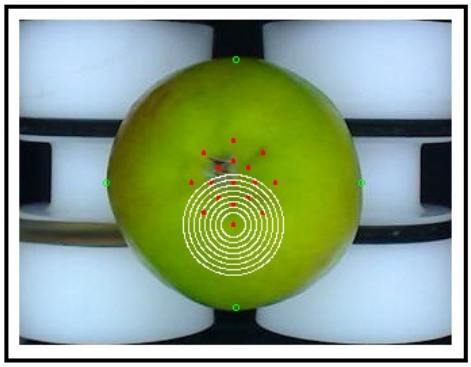

Figure 4.26 Concentric rings at origin 14

The concentric ring filter was performed at each of the new origins. Figure 4.26 shows the white rings about origin 14. The error and error^2 values between the fixed normal curve and the plotted concentric rings, as used in the previous filter, were recorded for each origin point for an image. From the error and error^2 data the origin with the highest error and error^2 was noted as the point where an island of interest may be present. In Figure 4.26 origin 1 had the greatest error, which coincided with a point close to the calyx in the image. This multi origin filter was run against the random library sample with a regression algorithm finding the best threshold to determine if an island of interest was present in the image.

Table 4-18 Test results from error calculations between dataset and fixed curve

Stem/Calyx/Defect	error	error^2
True Positive	543	542
False Negative	208	209
Total	751	751

The regression algorithm identified the best thresholds for identifying the presence of an island were (error > 445) and ($error^2$ > 13216). This multiple origin method produced better results than the single origin method (Table 4-17) with 4% more stem, calyx and defect islands being identified when compared with Table 4-18. The fixed line error, Equation 11, did not work as well as expected with the multiple origin filter because of the reduced outer radius of the concentric ring plot. This reduced radius resulted in a poor fitting intensity profile which did not always follow the expected shape of the apple surface.

To reduce the effect of the smaller outer concentric ring radius the gradient of a parabola fitted to the concentric ring dataset was examined as per the initial test described in Section 4.4.6. At each origin point the gradient of the fitted parabola and the error and $error^2$ between this fitted line and the data set was computed for the origin with the most negative gradient (the origin with the greatest dimple). The regression algorithm was used to identify the best thresholds for the new gradient and errors dataset. The regression algorithm identified the best cut-off threshold for an island of interest to be present at (Gradient < -0.09).

Table 4-19 Test results from error calculations between dataset and fitted curve

Stem/Calyx/Defect	Gradient
True Positive	637
False Negative	114
Total	751

The results, Table 4-19, from testing the gradient of the line fitted to the concentric rings dataset show an improvement of 12.5% over the fixed line test (Table 4-18). This is a large increase in filter performance over the other concentric ring filters used to find islands of interest, but it still less successful than the initial colour and Sobel method which identified 91.3% of the islands of interest whereas this only identifies 84.8% of the islands.

The error and error^2 for the line fitted to the concentric ring dataset was analysed for the origin recorded in the previous gradient test. The regression algorithm was used to determine if any useful information could be obtained from these line fit errors.

Table 4-20 Test results from error calculations between dataset and fitted curve

Stem/Calyx/Defect	error	error^2
True Positive	521	532
False Negative	230	219
Total	751	751

The regression algorithm identified thresholds of (error > 50) and (error^2 > 205) as the cut-off for an island of interest being present. The errors did not hold much useful information with at best 70.8% of the islands correctly identified by the error^2 in Table 4-20 which is worse than the results obtained from the gradient threshold. The errors did not produce any useful information because the smaller concentric ring radius, used in the multiple origin filter as opposed to the single origin filter, was more susceptible to noise due to the reduced number of pixels sampled and an island occurring near the edge of the rings affected the intensity profile of the plot more significantly. This was not a problem with the previous single origin concentric ring filter because if an island occurred near the edge it was diluted by the large number of points on the large outside rings.

4.4.8. Multiple origin asterisk filter

As a progression from the concentric rings, a star shaped filter was created that looked at the intensity profiles of lines that intersected the origin creating the star shape. The star resembles an asterisk with the three midpoints of the lines intersecting at the origin, resulting in a six point star shown in Figure 4.27. The intensity profile across the lines 1-3 (Figure 4.27) can be individually and collectively examined to determine if there is an island of interest present near the origin. The radius of the circle that would encompass the star was set as 1/3

the apple radius with the multiple origin r2, Figure 4.27, remaining at 1/3 apple radius. This (as in the previous concentric ring test 4.4.7) allowed a large portion of the apple to be examined while ensuring no background image would be processed (because the furthest point examined was located at a distance equal to 2/3 of the apple radius).

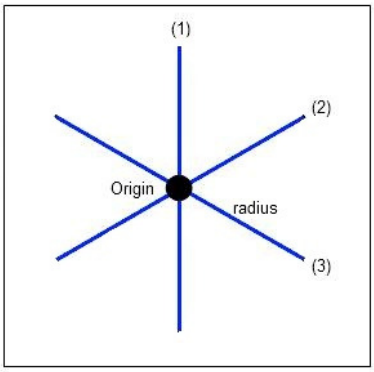

Figure 4.27 Dissecting lines of asterisk filter

The intensity data sets from the three lines produced by the asterisk (Figure 4.27) are averaged together to create a single intensity profile that can then be analysed to determine if an object of interest is present. This process is performed at each of the points in the multiple origin system, Figure 4.25, to ensure the most accurate results. The origin with the greatest gradient is identified as the point in that image where an object of interest is present.

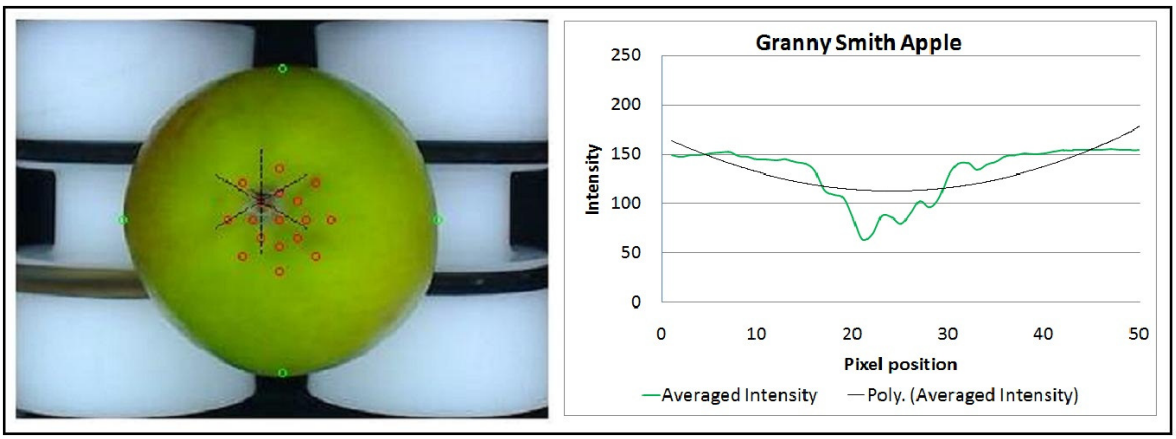

Figure 4.28 Asterisk Intensity profile for Granny Smith apple

The asterisk filter intensity profile of origin 9, Figure 4.28, shows the trend identified earlier where, if there is an object of interest present in a view, the light intensity dips causing a concave profile when a second order polynomial is fitted. The polynomial fitted to origin 9 in Figure 4.28, is identified as having the largest gradient of the 17 origins and thus is most likely to be the origin closest to a point of interest, which corresponds to the calyx visible in the image.

This asterisk filter was run against the random sample population library with the regression algorithm being used to determine the best threshold to identify if a stem, calyx or defect is present. The regression algorithm was set up to minimize the number of areas being misidentified by referencing the Stem/Calyx reference table to ensure an object of interest was present in a view when the threshold identified one as being present. The regression algorithm manipulated the threshold for the gradient of the line fitted to the averaged intensity dataset of a view. The more positive the gradient the more likely there was an object of interest present in that image.

Table 4-21 Test results for Asterisk search using gradient of fitted polynomial

Stem/Calyx/Defect	Gradient
True Positive	511
False Negative	240
Total	751

The regression algorithm identified a gradient threshold of 0.0638. With this threshold 511 stems, calyxes and defects were correctly identified (Table 4-21). This remains a lower accuracy than the colour and Sobel filter with 68% of stems, calyxes and defects correctly identified versus 91.3%.

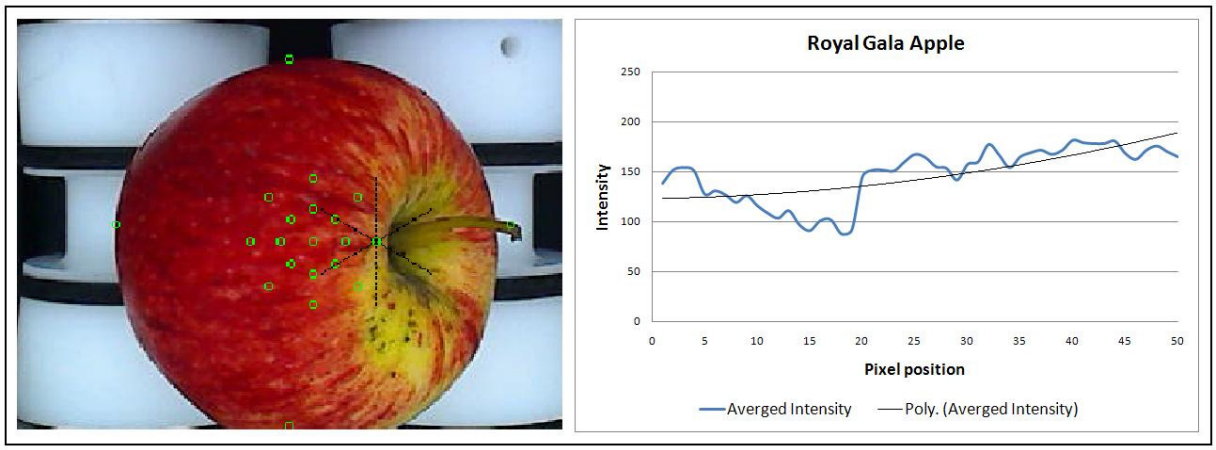

Figure 4.29 Asterisk search near edge of search boundary

Due to the shape of the asterisk filter, if an object of interest fell on the edge of the area being searched, it influenced the gradient of the fitted polynomial resulting in false positives the regression algorithm was trying to reduce. Figure 4.29 illustrates a stem on the edge of the search boundary. The intensity profile identifies the presence of an object while the Stem/Calyx reference table would say nothing is present because the stem is outside its search boundary. To reduce the effect of these outliers, the B component of the fitted polynomial, Equation (25), can be used to identify the centre axis of the fitted polynomial line. If the object of interest lies in the middle of the asterisk, similar to the calyx in Figure 4.28, the B component is near the midpoint of the asterisk while if the object of interest lies near the edge of the search boundary, similar to the stem in Figure 4.29, the B component of the fitted polynomial is off to one side. The midpoint for the asterisk search is pixel position 25. In Figure 4.28 (B = 23) while in Figure 4.29 (B = -4) resulting in a distinct difference between when the object of interest lies in the middle of the asterisk as opposed to its edge.

Using this discrimination the regression algorithm was run again allowing only asterisks with a centre axis of $15 < B < 35$ to be used to identify whether a stem, calyx or defect was present in the view.

Table 4-22 Test results for Asterisk search using gradient and position of fitted polynomial

Stem/Calyx/Defect	Gradient and Position
True Positive	606
False Negative	145
Total	751

The regression algorithm identified a new gradient threshold of 0.0314. The new filter correctly identified 80.7% of the stems, calyxes and defects. This was a large increase (12.6%) in the accuracy of the asterisk filter simply by including a threshold for the B component of the polynomial. This increase in filter performance resulted from an increase in accuracy of finding the correct multi point origin to examine rather than just the one with the greatest gradient where the centre of the object of interest may not have been in the centre of the asterisk search. Even with the increase in asterisk filter performance the colour and Sobel filter was still superior at finding any islands of interest that could be stems, calyxes or defects.

5. Discussion

In this chapter, the effectiveness of each of the filters designed to identify stems, calyxes and defects on the surface of the apples is discussed. The computational expense of each filter is stated and a comparison is made between commercial viability of each filter. The results of the work described in this thesis are then compared to work done by others.

5.1. Comparison of the filters

Of the six different searches proposed to identify stems, calyxes and defects on the surface of the apples the Sobel filter with Islanding routine (Section 4.4.1 and Section 4.4.2), proved the most accurate. In the sample population of 2,155 apple views this filter combination correctly identified 91.3% (686/751) of possible islands present along with a 5% (72/1404) false detection of islands of interest when no island was present. Of the total sample views 93.6% (2018/2155) views were correctly classified as either having or not having a stem, calyx or defect present. The false detections occurred when there was a sudden change in colour on bi-coloured apples or an object of interest was identified on the boundary of the search area but was not identified in the reference table because it was a partial view of the defect. A large portion of the unidentified stems, calyxes and defects were not identified because there was a gradual change in colour between the object of interest and the background so the Sobel filter was unable to identify these. After the threshold optimization, the Sobel filter with Islanding routine took 35 seconds to process the sample library of 2,155 images, i.e., an average of 16ms per image. The time to process each image was heavily dependent on the size of the search area and the number of pixels the islanding routine needed to process. An image with no identified pixels took less than 6ms to process while one that reached the island maximum population could take up to 42ms to be processed. Both of these are faster than the refresh rate of the cameras (66ms).

In an attempt to classify the areas of interest identified by the Sobel filter with Islanding routine as a stem, calyx or defect, the shape of each island was analysed using an Island characteristics filter (Section 4.4.3). The filter showed that while there was a discernable difference between the shape characteristics of stems, calyxes and defects during initial testing on specifically selected fruit this difference did not persist when the complete sample population was tested. There was a large overlap in characteristics thresholds between the stems, calyxes and different types of defects. This allowed islands to fall under more than one classification. In order to achieve a total classification accuracy of 97.4% over the sample population, 95.8% of islands fitted multiple classifications (Table 4-11). Only 1.6% of islands were identified correctly and exclusively as a stem, calyx or defect. The calculation of the island characteristics and comparison with the thresholds was again dependent on the island population. Calculation of the island characteristics took an average of 5ms per image on top of the time taken by the Sobel filter with Islanding routine. Larger islands took up to 11ms to process while smaller islands, such as those arising from pulled stems, took on average 3ms.

Promising results were achieved with a filter designed to identify whether the island of interest identified by the Sobel filter with Islanding routine occurred in a concave area of the apple. This was done by mapping the variation in geometric reflectance, about the island COG, due to the spherical nature of an apple (refer to the Concentric ring filter for colour graduations, Section 4.4.5). Using a regression algorithm to identify the optimal thresholds, 80.8% of stems and calyxes were correctly identified as being in a concave area while 81.1% of defects were correctly identified as being on the normal surface of the apple (Table 4-13). The success rate remained low because of the large variation in shape of apples. Some of the fruit had pronounced concave areas about the stem and calyx while others were flat. The filter relied on the shape being pronounced in order to distinguish between concave and convex areas similar to the samples shown during initial testing of the idea (Figure 4.15 and Figure 4.16). Had this method achieved a higher success rates it would have been an acceptable method to distinguish between stem/calyxes and defects. Calculation of the

concentric rings filter took an average of 12ms per image but this can be sped up by reducing the number of pixels being sampled. This computational period was in addition to the time taken by the Sobel filter with Islanding routine and the Island characteristics filter.

As a progression from the previous method, the variation in geometric reflectance due to the spherical nature of an apple was mapped about the apple's COG (Section 4.4.6). This removed the reliance on an island being found. Only 75% (568/751) stems, calyxes and defects were identified (Table 4-15) compared to an average of 91.3% (686/751) using the Sobel filter with Islanding routine. Because the variation in reflectance was mapped from the centre of the apple, an object of interest near the edge of the search area was not identified. The sample points were diluted by the large number of sample points recorded at the larger search radii of the outer rings. This method of identifying areas of interest was very fast with an average total processing time of 11ms per image but its poor accuracy makes it commercially unviable.

A significant improvement in the accuracy with which the variation in geometric reflectance is mapped was achieved by mapping it about several nodes (Section 4.4.7) at varying radii from the centre of the apple (Figure 4.25) as opposed to mapping from the centre of the apple only. The results from each point were then compared to identify whether an object of interest was present and where it was located. This greatly improved the accuracy with 84.7% (637/751) of stems, calyxes and defects being identified (Table 4-19) compared with 75% (568/751) using the previous method. Although this was a significant improvement it was still lower than that achieved by the Sobel filter with islanding routine. This filter was slower than the previous method because of the multiple nodes at which the concentric rings are processed. However, because the rings are of a smaller diameter, the sample size is smaller and the total processing time remains relatively fast at an average of 21ms per image.

The mapping of the variation in geometric reflectance was then altered to search three lines that intersected the origin nodes (Figure 4.25) creating the star shape (Figure 4.27) as opposed to concentric ring. This method (Section 4.4.8) had a somewhat lower accuracy with only 80.7% (606/751) of stems, calyxes and defects correctly identified (Table 4-22) after the search had been optimized. Compared with the concentric rings method, this method was more susceptible to localised noise because of the lower number of sample points. It was a fast filter taking an average of 17ms per image. Table 5-1 summarises the performance of the filters.

Table 5-1 Performance of filters

Filter type	Total processing time per image	Accuracy (%)*[2]
Sobel filter with Islanding routine	16 ms	91.3% of 751 images with blemishes
		93.6% of all 2,155 images
Sobel filter plus island characteristics filter	21 ms	97.4
Sobel filter plus island characteristics filter and Concentric ring filter for colour graduations	28 ms	81.0
Concentric ring filter for colour graduations	11 ms	75.0
Multiple origin concentric ring filter	21 ms	84.7
Multiple origin asterisk filter	17 ms	80.7

[2] The percentage of images in which all blemishes (stem, calyx or defect) were recognized without distinguishing amongst the blemish type

5.2. Comparison with other work

Table 5-2 summarises the main findings of published work in fruit inspection. A comparison between the inspection results reported here and the published literature has to take into account factors such as the size of the test sample, the time taken for inspection and the suitability of the method for commercial application.

Geoola et al. (1994) used spectrophotometric analysis with pixel thresholding to identify 91% of bruises on Golden Delicious apples while the stems and calyxes were hidden from view. The processing time is not stated but spectrophotometry is not rapid enough for commercial packhouse use. The research was limited to apples with a uniform skin colour.

Blasco et al. (2003) used colour segmentation to identify 95% of the stems and 86% of other defects on Golden Delicious apples, which have a homogenous skin colour. However, the method was tested on a very small image library (100 images of which 87 were of stems and only 13 were of other defects). Also, the segmentation method is very computationally intense (the processing time is not stated) and would fail on bi-coloured apples.

Using NIR pixel thresholding, image segmentation, neural networks and principal component analysis on a sample of only 200 images, Bennedsen and Perterson (2005) identified between 77% and 91% of defects with results varying across the eight different varieties of apples tested. In all images, the stem and calyx were hidden from view. The NIR spectrum was used to remove the colour textures from apples that were bi-coloured but even so, the method was most accurate for apples with uniform skin colour. The processing time was not stated but the authors state that this was a study to demonstrate feasibility of the method and that “no attempts were made to address or optimise the performance as far as speed was concerned.”

Unay and Gosselin (2006) reported a defect detection accuracy of up to 85% in bi-coloured Jonagold apples using image segmentation and a variety of unsupervised and supervised classifiers. The methods are described as “computationally very expensive” and would not be commercially feasible. The stems and calyxes were excluded as defects through templates matching. When low-resolution images were used, the processing speed increased but the accuracy of defect detection decreased.

When analysing citrus fruit using threshold segmentation in the visible, near infrared and ultra violet spectrums Blasco et al. (2007a) identified 63.4% of defects in the visible colour range, 79.5% of defects in the ultraviolet range and 92.9% in the near infrared range. No distinction between different defects could be made by the unsupervised classification algorithm. The method cannot be used on bi-coloured fruit.

Only Blasco et al. (2007b) and Lee et al. (2008) reported image processing times. Both were slower than the Sobel and Island routines reported in this work (Table 5-2) with Blasco et al’s citrus fruit segmentation routine being 250 times slower for a similar sized fruit and Lee et al’s date filter being 3 times slower for a significantly smaller fruit and monochrome image. Neither of these methods can be used on bi-coloured fruit.

Table 5-2 Summary of published work on fruit inspection

Author(s)	Fruit	Sample size	Detection Method	Accuracy (%)	Processing time per image (ms)	Comments
Geoola et al. (1994)	Apples (Golden Delicious)	300	Pixel thresholding	88.4% - 93.7% bruising detected	NS	Spectrophotometric analysis using NIR. Limited to fruit with uniform skin colour and hidden stems and calyxes.
Blasco et al. (2003)	Apples (Golden Delicious)	100	Colour segmentation and Bayesian nonlinear discrimination	93.4% of stems and 86% of the defects detected	NS	87 stem images and 13 non-stem images. Limited to fruit with uniform colour. Mechanically limited to four fruit per second
Bennedsen et al. (2005)	Apples (8 different varieties)	200	Image segmentation, neural networks and principal component analysis	77-91.6% defects detected	NS	Hidden stem and calyx in the image library. Highest accuracy for apples with uniform skin colour.
Unay et al. (2006)	Apples (Jonagold)	246	Image segmentation, pixel thresholding, 3 unsupervised classifiers and 10 supervised classifiers	<85% defects detected	NS	Stems and calyx regions removed from the images.
Xing et al. (2007)	Apples (3 different varieties)	329	Calculated softening ratio	95% Bruises detected	NS	No stem or calyx detection and assumes a large part of the fruit it healthy

Blasco et al. (2007a)	Oranges and mandarins	>2,294	Threshold segmentation with unsupervised classification	50% - 95.4% defects detected	NS	Uses UV, NIR and VIS light spectrums. Cannot distinguish between stems and defects. Limited to fruit with uniform skin colour and large difference in colour between fruit and defect. Upper limit of processing time: 30 minutes per image
Blasco et al. (2007b)	Citrus	2,132	Image segmentation	94% defects detected	4,000	Limited to fruit with uniform skin colour
Lee et al. (2008)	Dates	NS	Pixel thresholding and Sobel filter	74% -79% Grading accuracy	50	Uses NIR to identify skin delaminating
Balasundaram et al. (2009)	Citrus	180	Averaged Spectral reflectance	71%-100% Defect classification		Method suitable only for citrus
Caulton (2011)	Apples (4 different varieties)	2,155	Sobel filter with Islanding routine	93.6% of stems, calyxes and defects	16	Uses colour camera to identify stems, calyxes and defects.

NS: not stated

6. Conclusion

Several methods to inspect randomly orientated apples for defects were established and tested on a large sample of apples of four different varieties with varying success.

The high definition USB colour cameras proved a speedy and cheap method to capture images of the fruit as it was tumbled on a bi-cone roller arrangement. Although limited to 15 frames per second by the USB2.0 connection, the increased pixel resolution of these cameras allowed even small defects to be identified and more than one apple's surface to be captured in each frame.

Of the large sample population (2,155 images), 93.6% of the stems, calyxes and defects were correctly identified by a Sobel and islanding filter. Since the filter was not adjusted during testing to compensate for the different varieties of apples in the sample, and the speed with which it processed images, it can be concluded that the filter is suitable for an industrial application where multiple apple varieties are being rapidly processed. A filter to classify the shape of the identified islands could not reliably distinguish between a stem/calyx and a defect because of the large variation in their shape.

In an attempt to distinguish between the stems, calyxes and defects, the geometric reflectance of the apples surface was mapped with moderate success. This filter correctly identified 80.8% of stems and calyxes as being in a concave area and 81.1% of defects as being on the normal surface of the apple. This method showed promising results and with further work the accuracy could be increased to a level that would be accepted in the industry.

The filters reported in this work proved to be significantly faster than the methods of other researchers, many of whom were testing the feasibility of a method rather than its commercial viability. Some of these researchers got very good inspection accuracy on uniform-coloured fruit. The inspection accuracy achieved here was very high for both uniform- and bi-coloured apples.

7. References

Aleixos, N., J. Blasco, F. Navarron, E. Molto (2002). "Multispectral inspection of citrus in real-time using machine vision and digital signal processors." Computers and Electronics in Agriculture **33**(2): 121-137.

Balasundaram, D., T. F. Burks, D.M. Bulanon, T. Schubert, W.S. Lee (2009). "Spectral reflectance characteristics of citrus canker and other peel conditions of grapefruit." Postharvest Biology and Technology **51**(2): 220-226.

Bennedsen, B. S. and D. L. Peterson (2005). "Performance of a system for apple surface defect identification in near-infrared images." Biosystems Engineering **90**(4): 419-431.

Blasco, J., N. Aleixos, E. Molto (2003). "Machine Vision System for Automatic Quality Grading of Fruit." Biosystems Engineering **85**(4): 415-423.

Blasco, J., N. Aleixos, J. Gomez, E. Mollto (2007a). "Citrus sorting by identification of the most common defects using multispectral computer vision." Journal of Food Engineering **83**(3): 384-393.

Blasco, J., N. Aleixos, E. Molto (2007b). "Computer vision detection of peel defects in citrus by means of a region oriented segmentation algorithm." Journal of Food Engineering **81**(3): 535-543.

Blasco, J., N. Aleixos,S. Cubero, J. Gomez-Sanchis, E. Molto (2009a). "Automatic sorting of satsuma, (Citrus unshiu) segments using computer vision and morphological features." Computers and Electronics in Agriculture **66**(1): 1-8.

Blasco, J., S. Cubero, J. Gomez-Sanchis, P. Mira, E. Molto (2009b). "Development of a machine for the automatic sorting of pomegranate (Punica granatum) arils based on computer vision." Journal of Food Engineering **90**(1): 27-34.

Chayaprasert, W. and R. Stroshine (2005). "Rapid sensing of internal browning in whole apples using a low-cost, low-field proton magnetic resonance sensor." Postharvest Biology and Technology **36**(3): 291-301.

Davenel, A., C. Guizard, T. Labarre, F. Sevila (1988). "Automatic detection of surface defects on fruit by using a vision system." Journal of Agricultural Engineering Research **41**(1): 1-9.

Geoola, F., F. Geoola, U.M. Peiper (1994). "A Spectrophotometric Method for Detecting Surface Bruises on "Golden Delicious" Apples." Journal of Agricultural Engineering Research **58**(1): 47-51.

Gomez-Sanchis, J., E. Molto, G. Camps-Valls, L. Gomez-Chova, N. Aleixos, J. Blasco (2008). "Automatic correction of the effects of the light source on spherical objects. An application to the analysis of hyperspectral images of citrus fruits." Journal of Food Engineering **85**(2): 191-200.

Kader, A.A., (1992). "Postharvest Technology of Horticultural Crops". Second edition, Univ. Calif., Div. of Agr. and Nat. Resources, Publ. 3311, pp. 296.

Kleynen, O., V. Leemans, M.-F. Destain (2005). "Development of a multi-spectral vision system for the detection of defects on apples." Journal of Food Engineering **69**(1): 41-49.

Lee, D.-J., R. Schoenberger, J. Archibald, S. McCollum (2008). "Development of a machine vision system for automatic date grading using digital reflective near-infrared imaging." Journal of Food Engineering **86**(3): 388-398.

Leemans, V., H. Magein, M.F. Destain (1998). "Defects segmentation on 'Golden Delicious' apples by using colour machine vision." Computers and Electronics in Agriculture **20**(2): 117-130.

Lefcourt, A. M., P. Narayanan, U. Tasch, M.S. Kim, D. Reese, R. Rostamian, Y.M. Lo (2009). "Orienting apples for imaging using their inertial properties and random apple loading." Biosystems Engineering **104**(1): 64-71.

Li, Q., Wang, M., & Gu, W. (2002) "Computer vision based system for apples surface defect detection." Computers and Electronics in Agriculture, 36(2), 215-223.

Lloyd, E. L. (1983). "An O(n log m) algorithm for the Josephus Problem." Journal of Algorithms **4**(3): 262-270.

Mehl, P. M., Y. R. Chen, M.S. Kim, D.E. Chan (2004). "Development of hyperspectral imaging technique for the detection of apple surface defects and contaminations." Journal of Food Engineering **61**(1): 67-81.

Narayanan, P., A. A. Lefcourt, U. Tasch, R. Rostamiain (2008). "Orientation of apples using their internal properties" Transactions of the Asabe **51**(6): 2073-2081.

Njoroge, J. B., K. Ninomiya, N. Kondo, H. Toita (2002). Automated fruit grading system using image processing. Sice 2002: Proceedings of the 41st Sice Annual Conference, Vols 1-5**:** 1346-1351.

O'Rourke, A. D. (1994). The world apple market. Binghampton, NY: Food Products Press.

Penman, D. W. (2001). "Determination of stem and calyx location on apples using automatic visual inspection." Computers and Electronics in Agriculture **33**(1): 7-18.

Reese, D., A. M. Lefcourt, M. Kim, Y. Martin Lo (2009). "Using parabolic mirrors for complete imaging of apple surfaces." Bioresource Technology **100**(19): 4499-4506.

Studman, C. J. (2001). "Computers and electronics in postharvest technology - a review." Computers and Electronics in Agriculture **30**(1-3): 109-124.

Throop, J. A., D. J. Aneshansley, W.C Anger, D.L. Peterson (2005). "Quality evaluation of apples based on surface defects: development of an automated inspection system." Postharvest Biology and Technology **36**(3): 281-290.

Unay, D. and B. Gosselin (2006). "Automatic defect segmentation of 'Jonagold' apples on multi-spectral images: A comparative study." Postharvest Biology and Technology **42**(3): 271-279.

Van Zeebroeck, M., V. Van linden, H. Ramon, J. De Baerdemaeker, B.M. Nicolai, E. Tijskens (2007). "Impact damage of apples during transport and handling." Postharvest Biology and Technology **45**(2): 157-167.

Xing, J. and J. De Baerdemaeker (2007). "Fresh bruise detection by predicting softening index of apple tissue using VIS/NIR spectroscopy." Postharvest Biology and Technology **45**(2): 176-183.

Yang, C. C. and M. Marefat (1994). "Object oriented concepts and mechanisms for feature-based computer integrated inspection." Advances in Engineering Software **20**(2-3): 157-179.

Zion, B., S. M. Kim, M.J. McCarthy, P. Chen (1997). "Detection of pits in olives under motion by nuclear magnetic resonance." Journal of the Science of Food and Agriculture **75**(4): 496-502.