An Investigation of a Pattern Recognition System to Analyse and Classify Dried Fruit

Karen Henry

Submitted to the Faculty of Engineering, University of Cape Town, in partial fulfilment of the requirements for the degree Master of Applied Science in Engineering.
Cape Town,
April 1996

Declaration

I declare that this dissertation is my own work. It is being submitted for the degree of Master of Applied Science in Engineering at the University of Cape Town. It has not been submitted before for any degree or examination at this or any other university.

Signed by candidate

(Signature of Candidate)

Acknowledgements

I would like to thank my supervisor, Professor Gerhard de Jager, for his enthusiasm and guidance throughout this project.

I would also like to thank:

- Mr Peter Winm and his staff of South African Dried Fruit, for the effort and time they spent both with me and setting up the large test sets of dried fruit, without which most of this thesis would not have been possible.
- All of my colleagues, particularly Trevor Bartleet, Alan Bub and Fred Nicolls, for their advice and assistance, and for creating a great working environment.
- My family and Gavin Ashwell for their continued support and encouragement.

Abstract

Both the declining cost and increasing capabilities of specialised computer hardware for image processing have enabled computer vision systems to become a viable alternative to human visual inspection in industrial applications. In this thesis a vision system that will analyse and classify dried fruit is investigated. In human visual inspection of dried fruit, the colour of the fruit is often the main determinant of its grade; in specific cases the presence of blemishes and geometrical fault are also incorporated in order to determine the fruit grade.

A colour model that would successfully represent the colour variations within dried fruit grades, was investigated. The selected colour feature space formed the basis of a classification system which automatically allocated a sample unit of dried fruit to one specific grade. Various classification methods were investigated, and that which suited the system data and parameters was selected and evaluated using test sets of three types of dried fruit.

In order to successfully grade dried fruit, a number of additional problems had to be catered for: the red/brown coloured central core area of dried peaches had to be removed from the colour analysis, and Black blemishes upon dried pears had to be isolated and sized in order to supplement the colour classifier in the final classification of the pear. The core area of a dried peach was isolated using the Morphological Top-Hat transform, and Black blemishes upon pears were isolated using colour histogram thresholding techniques.

The test results indicated that although colour classification was the major determinant in the grading of dried fruit, other characteristics of the fruit had to be incorporated to achieve successful final classification results; these characteristics may be different for different types of dried fruit, but in the case of dried apricots, dried peaches and dried pears, they include the: peach core area removal, fruit geometry validation, and dried pear blemish isolation and sizing.

Contents

CONTENTS

List of Figures

LIST OF FIGURES

Chapter 1

Introduction

South Africa has a very successful and lucrative fruit industry. Some of the fruit that is harvested, is sulphur dried and sold to the dried fruit market. The quality control of both the fresh and dried fruit is very important. The quality grading of dried fruit is done either manually or by automation. Machine vision systems are available which grade both fresh and dried fruit by colour. Imported machine vision systems are expensive [34, 50]. They are very successful in fresh fruit sorting, but as found by South African Dried Fruit [34], are only partially effective in sorting dried fruit. A locally designed [9], and implemented machine vision system exists, for the colour sorting of fresh fruit. Because of mechanical considerations, this machine can only sort fresh fruit: the conveying belt includes large rollers to move the fruit, and weighing mechanisms, both of which are unsuitable for dried fruit. This machine vision grading system grades fruit by colour alone, as do the imported machine vision systems that were investigated. A locally developed machine vision system was required that would grade dried fruit.

In-depth discussions with [34], [50], [26], and careful observation of their respective fresh or dried fruit machine vision grading systems, indicated the need for a machine vision system that would cater for the following issues if it was to sort dried fruit successfully :

1. A colour classification system that could differentiate between the image colours of different grades of dried fruit. If each dried fruit that was examined was described according to its colour on a particular colour scale, then different dried fruit could be compared and consequently classified into a specific grade according to their value on the colour scale. This colour classification scheme would be the basis of a final dried fruit classification system that would tabulate and grade each dried fruit independently, into predefined classes or grades, according to its colour scale values.

2. In the South African dried fruit industry, the main dried fruit types include the dried

apricot, peach and pear. In the majority of cases dried fruit is sorted into different grades according to flesh colour. The machine vision systems that have been developed, parallel this system; fruit are graded according to their colour alone. The machine vision system of grading purely by colour has, however, been unsuccessful in dried peaches. In the case of dried apricots and peaches, before the fruit is sulphur dried, the pip of the fruit is removed. When the pip of a dried peach is removed, it leaves a core area that is much darker coloured than the remainder of the peach flesh; this problem does not exist with dried apricots. In the available machine vision systems, if the peach is graded according to its entire flesh colour, it is always graded as 'REJECT', because the core area is almost identical in colour to a peach blemish. A solution to this problem was required. The solution would require that any dried peach image would have to be divided into two parts; the dark coloured core area of a peach would have to be isolated from the remaining lighter coloured peach flesh. A process would have to be found that would accurately segment the core area in any dried peach image from the remaining peach flesh, in an acceptable time frame. This segmentation process would have to be performed on every single dried peach that was examined; the remaining non-core area of the flesh would then be used to classify and grade the peach according to its colour. The segmentation technique would have to make use of some of the properties of the peach image, for example the geometry, colour, or texture, in order to isolate the core area within the peach.

3. The grade specifications (i.e. the rules according to which a fruit is graded), although largely dependent on colour, include examination of two further areas; the fruit geometry, and in the specific case of dried pears, blemish area or size. Certain grades of dried pear are graded entirely according to the existence and size of black blemishes found on the flesh. Classification of these grades of dried pear require the recognition, isolation and sizing of black blemishes on the dried pears. This procedure needed to be automated. Both the dried fruit geometry and the blemish size are currently included in the final grading decision or classification of dried fruit. The necessity of automatic geometry validation needed to be investigated.

This thesis has been approached as, and constitutes a feasibility study on the most efficient way of automatically grading dried fruit. Fresh fruit were tested in parallel at each stage; good results were always achieved, ensuring a higher level of generality for the machine vision system if it was ever implemented. Although the final procedures that were selected for each problem were those that proved to be the fastest as well as most accurate, all testing was done using Khoros Image Processing [10] software; optimal times for the processes are therefore

not available and are beyond the scope of this investigation.

1.1 Format of this Dissertation

Chapter 2: The aim of this investigation was to design an accurate but efficient and time-effective machine vision grading system for dried fruit. In order to attain this, the problem of grading dried fruit using images captured using a CCD RGB camera, was examined and defined in the context of typical Pattern Recognition and Image Processing systems as described in [22], [27], [48], [54]. This process is detailed in Chapter 2. Pattern recognition systems have been grouped into three classes or 'concepts', and the qualifying characteristics of each, described. The dried fruit classification system is allocated to one of these general system types. Three system implementation methods are described and compared in order to highlight the most appropriate system design approach and theories for the planned dried fruit grading system.

Chapter 3: Visual colour is an accepted standard of measurement. Colour, therefore, is a major factor in the solution of visual, and hence automatic grading of dried fruit i.e. does this peach look bad to me? An introduction to human vision and colour is given in Chapter 3. A number of colour systems have been developed to 'calibrate' and reproduce colour, particularly in digital images. A number of these are described in this chapter, and the colour system that best suits the dried fruit grading problem, is investigated.

Chapter 4: Segmentation techniques to segment both (1) the fruit from the image background, and (2) black blemishes from within dried pears, were required. Thresholding techniques based on the colour histograms of the image proved very successful. Five thresholding techniques that analyse the image using the image histogram are described in detail in Chapter 4. The application of this thresholding technique to the two problems is described in Section 4.4 and 4.5.

Chapter 5: It was important that the dark core area of a dried peach was removed before the peach flesh could be classified according to a colour scheme. In order to segment the peach into core and non-core areas, the geometry of the peach image was considered. A geometric investigation of an image in Image Processing requires the theory and operations of Morphology. The basic theory and concepts of Morphology are introduced in Chapter 5, with special reference to the morphological operations which support the Top-Hat Transform operation. The application of the Top-Hat Transform to isolate the core area of a dried peach, is detailed at the end of Chapter 5.

Chapter 6: In order to examine and group each dried fruit according to its colour, certain classification techniques must be applied to the available information. Select classification

techniques are detailed and discussed in Chapter 6. Several of these techniques were investigated, and the optimum technique was selected.

Chapter 7: The test results for the colour classification system are detailed in this chapter. Chapter 7 is divided into three parts: a section that covers the test results of classification of dried apricots, a section that covers the test results of classification of dried peaches, and a section that covers the test results of classification of dried pears. In the case of dried pears, the results for a system that would incorporate both blemish analysis and colour classification in the final dried fruit grading system are presented.

Chapter 8: This chapter concludes and summarises the investigation of a machine vision dried fruit grading system. This is done with the aid of procedural flow charts: one is done for dried apricots, one for dried peaches and one for dried pears. A number of recommendations are proposed.

Appendix This includes summary discussions on :

- Details on the experimental hardware configuration.
- Testing for the front or back of a dried peach using simple statistics.
- Finding the most appropriate colour for the image background.

Chapter 2

Design of a Pattern Recognition System

2.1 Introduction

The aim of this investigation was to design a Pattern Recognition system (also called a machine vision system) that would analyse and then grade dried fruit. In order to design a machine vision system, it is necessary to consider the following :

- the form of the input data to the system : the images of the dried fruit that serve as input to the system, were captured using an RGB camera
- the output required from the system: *which* class is the dried fruit in ?
- to analyse the image data and decide the type of processing required to construct a final solution.

Lim [21] classifies the type of image processing system, where the input is an image, but the output is some symbolic representation of the contents of the image, as *image understanding.* Other examples of image processing applications would include *image restoration*, *image enhancement*, and *image coding.* In this chapter the basic process of system analysis, as described by Tou [22] is detailed. The process is divided into three steps; these are described in Section 2.3.

There are a number of approaches available to recognise image patterns. These approaches depend essentially on how the patterns in the image are characterised. Three pattern recognition approaches or concepts are described in Section 2.4: the Membership-roster Concept, the Common-property Concept, and the Clustering Concept. The problem of analysing and

grading dried fruit is compared to these Concept descriptions, and is related to that class of pattern recognition system that best describes it. The methods that would be appropriate to implement the system types of Section 2.5 are divided into three groups: Heuristic methods, Mathematical methods and Syntactic methods. Although the Concepts and Methods described in Section 2.4 and 2.5 are found in most text on Pattern Recognition Systems, the descriptions used here are those of Fu [27]. Once the pattern recognition system that is being investigated has been categorised, albeit very broadly, into a system type with a specific solution method, then a more distinct approach to the entire system solution is available.

Before the System Design process is described, a short list of issues that are relevant to machine vision systems are listed in Section 2.2. This list covers the more obvious advantages of using a machine vision system, as well as a number of factors that may restrict the choice of a machine vision system.

2.2 Advantages of a Pattern Recognition System

The most obvious advantages of a Pattern Recognition (machine vision) system are :

- Speed
- Accuracy
- Thoroughness (i.e. every item, rather than a random sample can be checked)
- Reliability
- Cost-effective (machine can run for 24 hours)

Factors which must be taken into consideration are :

- Cost of equipment
- Necessity of a controlled environment
- The undesirable effect of varying lighting (illumination) on results. Variations in illumination can be controlled or eliminated in controlled environments.

2.3 Pattern Recognition System Design - The Process

The design of an automatic pattern recognition system generally involves several major areas. Tou [22] describes them broadly as the following:

The first problem, known as the *sensing* problem, is concerned with the representation of input data which can be measured from the objects to be recognised. Each measured quantity describes a characteristic of the pattern or object. These measurements can be arranged in the form of a *measurement* or pattern vector

$$x = (x_1, x_2, .., x_n)$$

where each element of the vector x_i is a specific measurement type and $(i = 1..n)$.

The *measured data* in this problem was the reflected Red, Green and Blue values measured by the CCD [1] RGB camera. Each image scene (i.e. of several dried fruit), was described by the Red, Green and Blue light reflected from the surface of the dried fruit.

The second problem in pattern recognition concerns the extraction of characteristic *features* or attributes from the received input data, and thereafter the reduction of the dimensionality of this *feature* vector. The features of a pattern class are the characterising attributes common to all patterns belonging to that class. The primary justification of *feature* selection and dimensionality reduction, stems from engineering considerations [48]. As the complexity of a Pattern Recognition Classifier and its hardware implementation grow rapidly with the number of dimensions of the pattern space, it is important to base decisions only on the essential discriminatory information, which is conveyed by the *features*. Features are supposed to be invariant, or rather, less sensitive with respect to the commonly encountered variations and distortions that affect the input data. If a complete set of discriminating features for each pattern class can be determined from the measured data, the recognition and classification of these patterns will be simple:- in practice this is not often possible.

In this problem the characteristic which best described the dried fruit, was colour; dried fruit are graded largely according to their flesh colour. Various colour coordinate systems exist for describing the colour in an image. Both *quantitative* and *perceptual* colour models are discussed in Chapter 3. The two most promising colour systems were examined. From these results a feature vector (i.e. a particular colour system) was selected. In order to reduce the dimensionality of the feature vector, each of the vector elements were examined in detail; those vector elements that could be dismissed without causing an unacceptable loss in image information, were removed from the feature vector.

The third problem in pattern recognition systems involves the determination of an optimum decision procedure. This is needed in the identification and classification process. After the observed data from patterns have been expressed in the form of *measurement vectors* (also

[1] The charge coupled device (CCD) is an example of a solid-state sensor element; a 2-Dim array of sensor elements is integrated on a chip. One sensor element is located spatially at each pixel location and senses the light intensity at each pixel. When a CCD array is exposed to light, charge packets proportional to the light intensity develop. These packets are moved to a storage array, from which the Intensity values are then read.

called *feature vectors*) in the *feature space*, the machine must decide to which pattern class these data belong.

If each dried fruit was described by a feature vector $[x_1, x_2, x_3]$, and two possible fruit grades existed: *good* and *bad*, then a method would be required to decide whether a fruit belonged to one grade or the other. This method would include a decision procedure that would allocate a fruit to either one grade or the other. In Chapter 6, a number of decision procedures relevant to the current classification problem description are outlined. These procedures were examined, and the optimum decision procedure was selected.

2.4 System Design Concepts

The design concepts for automatic pattern recognition are motivated by the ways in which pattern classes are characterised and defined. Several possibilities are the :

- Membership-roster concept
- Common-property concept
- Clustering concept

2.4.1 Pattern Recognition using a Membership-roster Concept

Characterisation of a pattern class by a roster of its members suggests automatic pattern recognition by *template matching*. The set of patterns for each class are stored in the system. When an unknown pattern is shown to the system, it is compared to each of the stored patterns. If the input pattern matches a stored pattern, then the input pattern becomes a member of that particular class. This approach, however, requires the input of near perfect pattern samples.

2.4.2 Pattern Recognition using a Common-property Concept

In this approach, a pattern class is characterised by properties or attributes common to that specific class. These common properties can be stored in the recognition system. When an unknown pattern is observed by the system, its features are extracted, possibly manipulated, and compared with the stored features. The recognition scheme will classify the new pattern as belonging to the pattern class with similar features. If all of the features of a class can be determined from sample patterns, the recognition process reduces to simple 'feature match-

ing'. As this is seldom the case, one needs to find those class features which are optimum in some sense. Tou [22] discusses several methods for feature selection.

2.4.3 Pattern Recognition using a Clustering Concept

When the patterns of a class are vectors whose components are real numbers, a pattern class can be characterised by its clustering properties in the pattern space. Points in the pattern space which cluster together are linked into one class. How the clusters are linked (i.e. chains of points, isolated groups of points) is decided by the system designer. The relative geometrical arrangement of the clusters will dictate the recognition scheme. If the classes are characterised by clusters which are far apart, simple recognition schemes may be used; if the clusters overlap, more sophisticated techniques are necessary to partition the pattern space.

2.4.4 Basic Pattern Recognition for Dried Fruit Grading

In a fruit classification system, the aim of the recognition system is to examine the digital image of a dried fruit and to decide which class (grade) that fruit belongs to. People generally judge visually whether a fruit is *good* or *bad*; when deciding which class to allocate the fruit to, a person would look at the colour of the fruit, compare that colour to the colour of '*good* fruit' which they have stored in their memory, and decide whether that sample fruit was, or was not, the same colour. A parallel pattern recognition system would make use of either the common-property, or, more quantitatively, the clustering concept. Both concepts require the use of characteristic features, although the Clustering Concept requires a numeric real-number representation of the feature. Features that best described *good* and *bad* fruit had to be found. If the pattern or feature vectors, which described the dried fruit, clustered into geometrically distant groups for *good* and *bad* fruit samples, clustering would be an ideal technique.

2.5 Implementation Methods for Selected Concepts

Tou grouped the methods used to implement the above pattern recognition concepts into three principal categories :

- Heuristic methods
- Mathematical methods
- Syntactic methods

Section 2.5: Implementation Methods for Selected Concepts

2.5.1 Heuristic Methods

This method is based on human experience and intuition, and depends to a large degree upon the capability of the system designer. The concepts of *Membership-roster* and *Common-property* are used. Procedures are developed for specialised recognition tasks. An example would be the problem of character recognition where the classification of a character could be based on the detection of features such as the number and sequence of particular strokes.

2.5.2 Mathematical Methods

The Mathematical approach is based on classification rules which are formulated and derived in a mathematical framework, making use of the common-property and clustering concepts. The mathematical approach can be divided into two categories : *deterministic* and *statistical.* The *deterministic* approach is based on a framework which does not employ the statistical properties of the pattern classes. The *statistical* approach is based on mathematical classification rules which are formulated and derived in a statistical framework; this classifier is generally based on the Bayes classification rule and its variations [27].

2.5.3 Syntactic(Linguistic) Methods

In this method patterns are characterised by their sub patterns or *primitive* elements. The approach of structuring primitives into a hierarchy in order to determine a pattern is known as Syntactic Pattern Recognition, and is analogous to the syntactic structure of languages where the pattern grammar consists of finite sets of elements called variables, primitives and productions. A detailed discussion of Syntactic Pattern Recognition methods can be found in [27].

2.5.4 Method selected for Dried Fruit Grading

Each digital image consisted of a single dried fruit upon a unicoloured background. Only one object was in the image:- the dried fruit. The feature that most obviously suggested the grade (i.e. good or bad) of the fruit was the colour. Clearly the use of *primitives* was not suited to this non-hierarchical feature space; the Syntactic method was therefore not suitable. Although a Heuristic method could have been devised, the problem did require a classification system that would predict the grade of each new dried fruit; the Mathematical method was therefore selected as the most appropriate of the three methods. A *statistical* Mathematical approach required an *a priori* knowledge of the relative volumes expected in each fruit class (i.e. 50% of the dried peaches of every run would be in the Choice grade). These proportions

were unknowns, but could be assumed equal for test purposes.

The aim of this chapter was to investigate different Pattern Recognition systems, and the type of problems to which they are generally applicable:- this precursive analysis was useful in guiding the system analysis procedure.

In the next chapter the problem of *feature extraction* is examined. Various colour systems are described and investigated in order to select the feature space which will most successfully describe dried fruit, and highlight the differences between both **dried fruit types** (i.e. Apricot versus Peach), and the individual **dried fruit grades** (i.e. Choice versus Reject) of each fruit type.

Section 2.5: Implementation Methods for Selected Concepts

Chapter 3

Colour Perception and Measurement

3.1 Introduction

One of the main objectives of the dried fruit grading system that was investigated, was to find a colour system that adequately represented the colour differences within grades of dried fruit. The coordinates of the colour system would be used to represent each fruit uniquely; the coordinates would constitute the *features* of the *feature space.*

In this chapter various methods of representing image colour in image processing and other electronic display systems, are described. Both *perceptual* and *quantitative* colour coordinate systems are available to specify the colour in an image. *Quantitative* specification of colour is based on physical measurements of the light wavelengths reflected from an object [56]. *Perceptual* colour specification, on the other hand, represents the range of colours that a human perceives; these type of colour systems are sometimes referred to as the 'artist's palette' [17]. The *Perceptual* colour models that are discussed in Section 3.4.2 include the HSV (Hue, Saturation, Value or Intensity) model and the HLS (Hue, Lightness, Saturation) model. The *Quantitative* colour representations detailed in Section 3.4.4 include the RGB (Red, Green, Blue), CMY (Cyan, Magenta, Yellow) and YIQ colour models. The Commission Internationale de l'Eclairage (CIE) International Colour System is introduced in Section 3.4.3. This System is the basis of the quantitative RGB, CMY and YIQ colour models. More expansive descriptions of the colour models introduced in this chapter can be found in [56], [9], and [17].

In the selection of a colour system, two issues were important :

Section 3.1: Introduction

1. The colour system should clearly differentiate between the colours of different grades of dried fruit, within the said colour feature space.

2. The overall processing time of any commercial machine vision system is very important. Any colour system that enhances the overall processing speed of the classification scheme should be considered seriously; conversely any colour systems that would lengthen this process, would have to be viewed with caution.

Because of the large number of colour systems that are available, it would be an arduous task to sufficiently investigate each colour system with respect to the proposed dried fruit colour classification scheme. Instead, the approach was 3-fold :

1. Available literature that had investigated colour models for colour segmentation, was examined.

2. Machine vision systems that are available which classify fruit (dried or fresh) by colour, were investigated.

3. A final analysis was done to compare the two most suitable colour models, derived from the above two steps, and to select one.

Ohta *et al* [57] documented an investigation of colour segmentation, comparing a broad range of colour models; they conducted a very thorough and comprehensive investigation of more than 100 colour features (including RGB, YIQ, HSI, normalised RGB, and La*b* systems), in order to find those colour features which achieved a colour segmentation comparable to that of the colour features X_1,X_2,X_3 obtained by the reputable Karhunen-Loeve transformation [53]. Three colour systems performed particularly well: the La*b*, the HSI, and the derived I_1,I_2,I_3 (described below) colour systems.

Only the $L^*u^*v^*$ and the $L^*a^*b^*$ models, as described by [16] are recommended by the Commission Internationale de l'Eclairage (CIE) for uniform colour difference evaluation. It was found that the La*b* model, although producing good results was too costly; from the computational point of view, the non-linear transformations incur far more cost than linear transformations. Uchiyama and Arbib [47] reached a similar conclusion on the La*b* model (or similarly the Lu*v* model) when investigating select colour models that would effectively segment colour image scenes. The subject of uniform colour spaces is discussed in [35], [16] and applied in [29] Ohta *et al* [57] concluded that the three orthogonal colour features $I_1 = \frac{(R+G+B)}{3}$, $I_2 = R - B$ and $I_3 = \frac{2G-R-B}{2}$ were important components for representing colour information. Ohta tested the success of using only two of these features, I_1, I_2 to segment the examples: the results were comparable to those achieved by the three features, except these two features failed to differentiate between pale yellow, yellow, golden yellow

and orange. Differentiation of these colours, however is vital in a dried fruit grading system. The I_3 feature is heavily weighted by the Green value: this result highlights the importance of the Green value in the successful recognition of orange/yellow colours, and the possibile success of the I_1 and I_3 features in describing the colour grades of dried fruit. Sun *et al* [58] implemented Ohta's I_1,I_2,I_3 colour features to successfully analyse liver tissue.

It was noted by Ohta *et al* that various kinds of colour features such as Intensity, Saturation, and Hue are calculated from the tristimuli values of Red(R), Green(G), and Blue(B) by using either linear or non-linear transformations. This reliance on the R,G,B values means however, that normalised RGB colour and Saturation (see Section 3.4.2 for calculation) values become unstable and meaningless when $R + G + B$ is very small; in this regard it should be noted that dried fruit blemishes have distinctly low values of Green and Blue colour. Miller [5] used normalised Red and Green values ($r = \frac{R}{R+G+B}$,$g = \frac{G}{R+G+B}$) to grade fresh peaches. Miller classified each peach according to the minimum Euclidean distance between the (r,g) coordinates of the peach and the standard (r,g) values designated to represent each peach class. The final fruit classifications according to this system were correct only 54% of the time, with 88% of classification results being within one colour reference (class) of the manual colour estimate. These bad results are likely to be the result of assuming Euclidean distance variation in the (r,g) variables, rather than statistical distance variation. (The Euclidean and statistical variations in the Red and Green values of both dried and fresh fruit were examined and are discussed in detail in Chapter 6.)

Machine vision systems that have been developed to analyse and classify fruit according to their colour include: the fruit sorter designed by Kay [9] and implemented by Van Zyl [50] which uses the HSI colour model; the KROMCO [26] fruit sorter, which uses the HSI model; a design for a fruit sorter by Guedalia [13], which makes use of the HSI model and the KROMA SORTER [43] which uses the RGB colour model.

The conclusions derived from the literature and relevant machine vision fruit sorting systems were that :

1. Ohta *et al*'s colour features I_1,I_2,I_3 based on a simple transformation of R,G,B values defined different colours well.

2. The HSI system, although requiring time-consuming matrix conversions, was favoured in many of the machine vision sorting systems, and is successful in commercial machines.

3. The La*b* system, although it shows favourable results, requires time-consuming matrix conversions from RGB space.

It was decided that further investigation of both the RGB and the HSI colour systems was necessary. Because the HSI model is a colour model which reflects those colours seen in

human vision, it was expected that this model would clearly differentiate between the colours in dried fruit grades. What had to be investigated was whether or not the RGB model would recognise the colour differences as effectively as the HSI model. For both the HSI and RGB colour systems, the *feature space* was examined with the intention of reducing the dimensionality. In this case it meant selecting only one or two of the three coordinates in either the HSI or RGB systems, to represent the colour of a dried fruit. This investigation is the subject of Section 3.5.

Human vision is not a simple process; what the human eye perceives is a result of physiological processes [56], as well as environmental factors. The effect of environmental factors on what we see, as described by Hunt [40], is illustrated by way of example in Section 3.2. Section 3.3 deals with human vision. Section 3.4 deals with the specification of colour: this includes both *quantitative* and *perceptual* colour models. Finally Section 3.5 details the investigation and comparison of the HSI and RGB colour models.

3.2 Human Vision

Consider a street in daylight. All of the light falling on the street comes from the sun. Since sunlight is a combination of all of the colours of the spectrum, the street scene is being illuminated by this mixture of wavelengths. Foliage contains a dye called chlorophyll which has the property of absorbing red and blue light but transmitting (reflecting) green light; hence when foliage is illuminated by daylight only the greenish components of the incident light are visible to the eye - we say the foliage looks green. If we return to the street after dark and find that it is lit by sodium lamps we shall find that the leaves now look black. The illuminant contains only yellow light which is entirely absorbed by the foliage i.e. there is no green light for the foliage to reflect. It is clear that both the quality and type of illuminant, as well as the nature of the object, contribute toward the colour we see.

Vision is a very powerful sense. It provides us with a remarkable amount of information about our surroundings. In order to understand the nature of colour vision, it is necessary to understand the basic physiology of human vision. This is introduced in the following section.

3.3 The Physiology of Human Colour Vision

The visual system consists of an eye that transforms light to neural signals, and the related parts of the brain that process the neural signals and extract necessary information. From a functional point of view, the eye is a device that gathers light and focuses it on its rear

surface, the retina [21]. The human retina contains two main types of light receptive cells known as *rods* and *cones*. Colour vision is simulated in the *cones* of the retina. The cones are most sensitive to Red, Green and Blue light. This is the quantitative physiological basis for representing a colour image with Red, Green and Blue monochrome images [21]. The cones are less sensitive to light than the rods. The colourless vision which occurs at very low levels of illumination, such as weak moonlight, depends on the bleaching of a photo-sensitive substance called 'visual purple' contained in the rods.

The following Section introduces the current methods and ways of specifying an observed colour.

3.4 The Specification of Colour

The analysis and description of colour can be viewed *subjectively* or *quantitatively*, but for the purpose of technical specifications and machine vision, consistent, tabulated colour models are required. In Section 3.4.1 the terminology for *subjective* measurement is described, together with two subjective colour models; the HSV and HLS models. In Section 3.4.3 the *quantitative* colour system of CIE, and the *gamuts* or colour models of RGB, YIQ, and CMY are described. All of the colour models described in this section are adopted from the colour model representations in [17]. These and other important colour systems and their transformations, are also detailed in [53].

3.4.1 Perceptual Specification

Subjectively, colours have three main attributes; *hue*, *saturation* and *lightness*. These are generally referred to as the *artist's palette*, as they are closely linked to a humans perception of colour measurement. The three attributes can be described as follows :

hue is the colour the object appears to be e.g. red or orange. The hue of an object can vary according to background, lighting and various other factors.

saturation denotes the purity of the colour. It is an attribute that relates the position on a scale which extends from a bright colour, through increasingly dull colours, to grey. The extent of colour saturation is dependent upon :

1. the general level of lighting intensity,
2. how directional the lighting is,
3. the type of surface of the object [40].

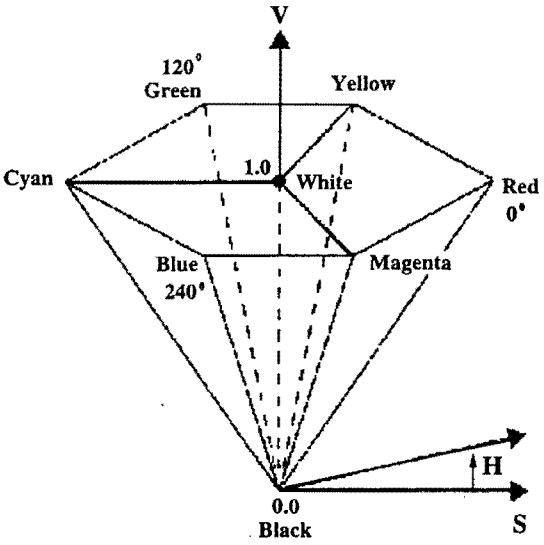

Figure 3.1: Single-hexcone HSV model.

lightness is a quantity based upon the relative brightness of the colour of an object(which does not emit light). The scale is measured in equal steps ranging between ideal Black and ideal White.

3.4.2 Perceptual Colour Models

HSV Colour Model

This model of Hue(H), Saturation(S) and Value(V) (also called brightness B, or Intensity I, in the identical HSB and HSI models), is not hardware based, but is related to an artist's tint, shade and tone. There is no single method of calculating HSV values; a number of methods are compared by Kay [9]. One of these methods is the method of Niblack [51] :

The Hue is set to zero in the direction of Blue, and the Intensity(I) is R=G=B. The vector v_1 is orthogonal to I in the plane of I and B, and the vector v_2 is I cross v_1. The transformation is :

$$\begin{pmatrix} I \\ v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{-1}{\sqrt{6}} & \frac{-1}{\sqrt{6}} & \frac{-2}{\sqrt{6}} \\ \frac{-1}{\sqrt{6}} & \frac{-2}{\sqrt{6}} & 0 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

The final calculation for H and S is :

$$H = tan^{-1}(\frac{v_2}{v_1})$$

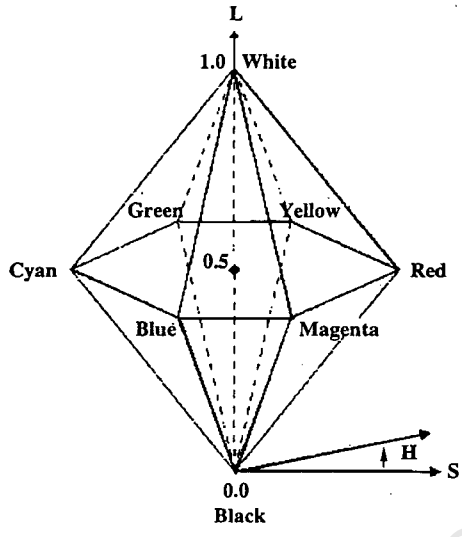

Figure 3.2: Double-hexcone HLS model.

$$S = \sqrt{{v_1}^2 + {v_2}^2} \tag{3.1}$$

where H ranges from 0 to 360^o.

This coordinate system contains three coordinates of Hue, Saturation and Value and is detailed in Figure 3.1. In this Figure the Hue is measured by the angle around the vertical axis V with Red at zero degrees. The Saturation (S) is a ratio ranging from zero on the V axis to 1 on the sides of the hexcone. A specific colour is pure when S = 1 and V = 1. The Value (V) on the vertical axis has a range of $[0, 1]$ where 0.0 denotes Black and 1 denotes White. The intermediate values of V denote the greys.

HLS Colour Model

The Hue(H), Lightness(L), Saturation(S) colour model is very similar to the HSV model, except that Lightness scales from L = 0, denoting Black, to L = 1, denoting White, on a double-hexcone subset of a cylindrical space. The Hue angle choices are not fixed i.e. Red is not always found at zero degrees. The most pure Hues are found at S = 1, L = 0.5.

A more objective, quantitative method of specifying colour, called *Colorimetry*, which is described extensively by Grand [56], is introduced in the next section.

3.4.3 Quantitative Specification - Colorimetry

Colorimetry is the science of quantitatively measuring colour [53]. A rigorous explanation of quantitative colour matching and the relevant mathematical approach, is found in [53]. The important terms in *colorimetry*, and their perceptual equivalents are listed below.

Perceptual Term	Colorimetry
Hue	Dominant wavelength
Saturation	Excitation purity
Lightness (reflecting objects)	Luminance
Lightness (self-luminous objects)	Luminance

A basic description of each term is as follows :

Dominant wavelength is the wavelength of the colour we see when viewing the light.

Excitation purity corresponds to the saturation or vividness of the colour. The excitation purity of a coloured light is the proportion of both the pure light of the dominant wavelength and the White light needed to define the final colour. A pure colour is therefore 100 percent saturated whereas non-pure colours have saturation between 0 and 100 percent. White light and Greys contain no colour of any dominant wavelength and are therefore zero percent saturated.

Luminance describes the amount or the intensity of the light.

The CIE Chromaticity Diagram

In 1931 the Commission Internationale de l'Eclairage (CIE) introduced a method whereby any colour could be defined using a mixture of three fixed primary colours, **X**, **Y** and **Z**. The amount or value of each primary **X**, **Y**, **Z** which are used to specify a colour **C**, are known as the *tristimulus values*; hence any colour **C** can be written $\mathbf{C} = X\mathbf{X} + Y\mathbf{Y} + Z\mathbf{Z}$. This was called the CIE Colorimetric System. The tristimulus values which were adopted by the CIE are represented graphically in Figure 3.3. The colour-matching functions of $\overline{x_\lambda}$, $\overline{y_\lambda}$, $\overline{z_\lambda}$, as illustrated by Foley [17] indicate the amount of each of the CIE primaries that is required to match the colour produced by one Watt of radiant power at a certain wavelength [17]. In other words the amounts of X, Y and Z primaries needed to match a colour with a spectral energy distribution $P(\lambda)$ (where $P(\lambda)$ is measured with a spectroradiometer) are :

$$X = k \int P(\lambda)\, \overline{x_\lambda}\, d\lambda$$

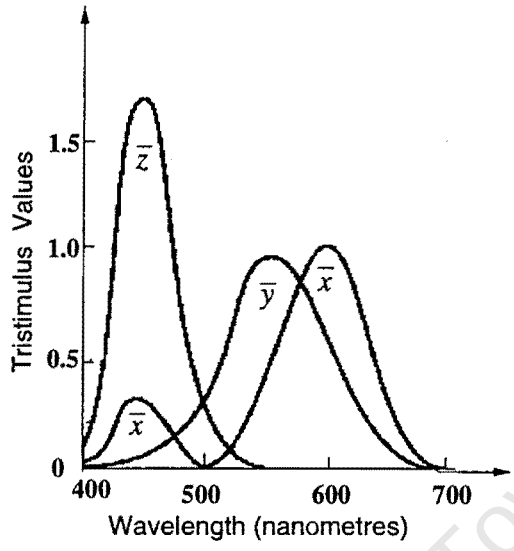

Figure 3.3: The colour matching functions $\overline{x}$, $\overline{y}$, and $\overline{z}$, for the CIE X, Y, Z primaries

$$Y = k \int P(\lambda)\, \overline{y_\lambda}\, d\lambda$$

$$Z = k \int P(\lambda)\, \overline{z_\lambda}\, d\lambda.$$

A set of *tristimulus* values that are frequently used are the R,G,B values which represent the intensity of the Red, Green and Blue light. These three colours are used because by proper combination, they can produce a wider range of colours than any other combination of three colours [21]. The units of these *tristimulus* values are such that the tristimulus values R,G,B are equal when matching an equal energy White throughout the visible spectrum [53].

Foley [17] derived a 2-dimensional (x, y) coordinate system from the CIE *tristimulus* co-ordinate system. The *chromaticity* coordinates (x, y) are defined by normalising against $(X + Y + Z)$, where $(X + Y + Z)$ can be thought of as the total amount of light energy.

$$x = \frac{X}{(X + Y + Z)} \tag{3.2}$$

$$y = \frac{Y}{(X + Y + Z)} \tag{3.3}$$

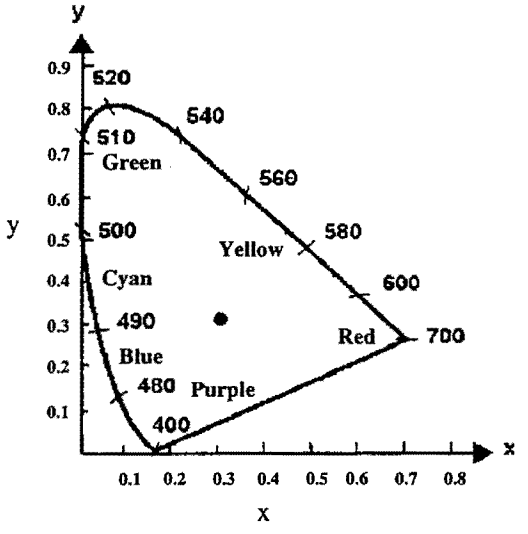

Figure 3.4: The CIE chromaticity diagram.

$$z = 1 - x - y \tag{3.4}$$

By plotting x and y for all visible colours, which is the range of 400 to 700 nanometres, we obtain the CIE chromaticity diagram described in Figure 3.4. The interior and the boundaries of this region represent all chromaticity values. Colours with the same chromaticity value but different luminances will map onto the same point in the diagram. The 100 percent pure colours are found on the curved part of the boundary. A standard White light approximating sunlight is defined by a light source illuminant C.

The chromaticity diagram in Figure 3.5 could be used graphically as follows :-

- When two colours are added together, the new colour lies somewhere along the straight line that connects the two colours being added. This result is defined by the 'Centre of Gravity Law' detailed by Hunt [40].
- Colour A can be thought of as a mixture of 'standard' White light (C) and a pure colour at B. B therefore defines the dominant wavelength.
- The ratio of length AC to length BC, expressed as a percentage, is the *excitation purity* of A. The closer A is to C, the less pure A is and the more White light it includes.

The RGB, YIQ and CMY models, which are subsets of the CIE colour space, are commonly used in raster graphics [17]. The purpose of these colour models is to allow specification of

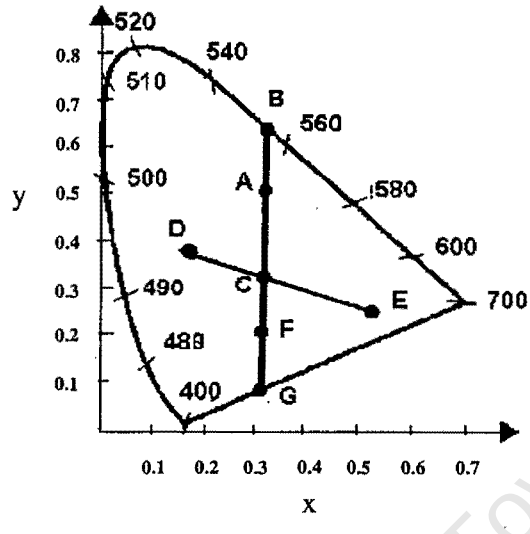

Figure 3.5: Colours on the chromaticity diagram.

colours within some colour *gamut*. The RGB, YIQ and CMY are three hardware-oriented colour models :-RGB is used with colour cathode ray-tubes (CRT) monitors, YIQ with Television colour systems and CMY with some colour printing devices. The matrix representations listed in this Section, for these models are derived from Foley [17].

3.4.4 Quantitative Colour Models

The RGB Colour Model

The Red(R), Green(G) and Blue(B) colour model uses a Cartesian coordinate system. As indicated graphically in Figure 3.6 the model describes a cube. The main diagonal of the cube, which has equal amounts of each primary R, G and B, represents the Grey levels. The colour Black is at position $(0,0,0)$ and White is at position $(1,1,1)$.

The colour gamut covered by the RGB model is defined by the chromaticities of a CRT monitor's phosphors. One can transfer between two different CRT colour gamuts by applying a transformation $m_{1\rightarrow 2}$:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = m \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

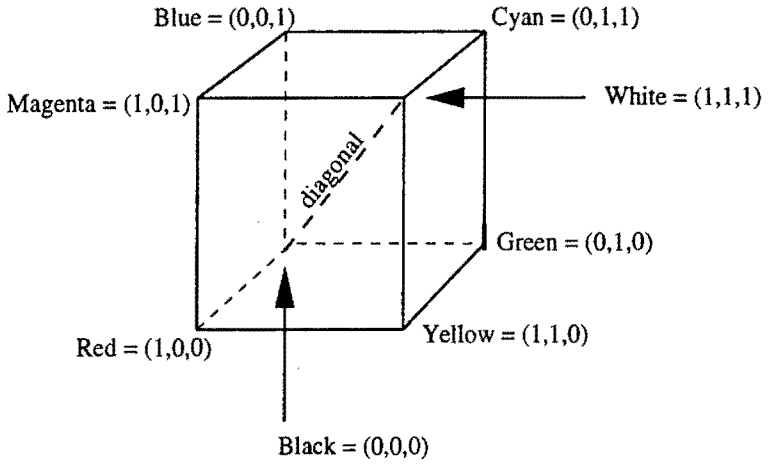

Figure 3.6: The RGB cube where Greys lie on the diagonal.

where

$$m = \begin{pmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{pmatrix}.$$

and X_r, X_g, X_b are the weights applied to the monitors RGB colours to find X in the new colour space.

The colour gamut for a typical colour monitor is clearly smaller than that available in CIE colour space.

CMY Colour Model

Nature often generates colour by filtering out, or subtracting, some wavelengths and reflecting others. The process of wavelength subtraction is accomplished by molecules called *pigments*, which absorb particular parts of the light spectrum :- this is called a *subtractive* colour system [21]. The three primary colours of a *subtractive* colour system are Cyan, Magenta and Yellow. Cyan, Magenta and Yellow are the complements of Red (R), Green (G) and Blue (B) respectively. A knowledge of the CMY colour model is important when dealing with devices that deposit coloured pigments onto paper (i.e. ink-jet plotters). When a surface is coated with Cyan ink, no Red light is reflected from the surface because the Cyan pigment absorbs, rather than reflects the Red light wavelengths and we say that Red is *subtracted* from the reflected White light. When all three colours Cyan, Magenta and Yellow are combined, the result is Black where the pigments together absorb all of the visible wavelengths.

The CMY model is the same as that for RGB except that White instead of Black is at the

origin.

$$\begin{pmatrix} C \\ M \\ Y \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} - \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

where $(1, 1, 1)$ is the RGB representation of white.

YIQ Colour Model

YIQ is used in commercial television broadcasting. This model is a recoding of the RGB model for transmission efficiency [17]. The Y component of YIQ is defined to be the same as the CIE Y primary of Luminance.

$$\begin{pmatrix} Y \\ I \\ Q \end{pmatrix} = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.528 & 0.311 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

3.5 Selection of a Colour Coordinate System

The colour limits and colour variations within a dried fruit class were known. What was required was a colour system that could clearly differentiate between the colours of different grades of dried fruit. In the final system a sample dried fruit would be compared to the known class colours and allocated to the class whose range of colours most closely resembled that of the sample fruit. In this section the colour systems of the HSI and the RGB models are examined and compared. Each of the variables in the two colour spaces were analysed independently to ensure that redundant features were discarded. Using an assortment of fresh and dried fruit types, the separation and deviation of the variables within *good* and *bad* grades were examined. For each of five different fruit types, the Euclidean distance d_E [1] between the two grades of *good* and *bad* fruit, the average value for each grade (average over three fruit samples), and the related standard deviation value are given in TABLE 1. These results are compared and the optimum feature space selected.

A number of the colour models discussed in this chapter are relevant to certain hardware applications : the CMY system is used for colour printing purposes and the PAL system is used for output to television screens. An investigation of relevant literature and of current machine

[1] The *Euclidean distance* is a real-valued distance function (d_E) [53] of two image points (j_1, k_1) and (j_2, k_2) where :

$$d_E = \sqrt{[(j_1 - j_2)^2 + (k_1 - k_2)^2]}$$

vision fruit sorting systems highlighted the quantitative RGB model and the perceptual HSV model (which will be referred to as HSI from here on, as this is the notation used in the Khoros system [10]), as the most appropriate colour models for a dried fruit grading system. Although the HSI colour system has no correlation with the standard CIE colour system [42], and requires a lengthy matrix transformation from the RGB source input (RGB camera), it describes the colour that the human eye observes; any system making use of the HSI model can be trusted to reasonably represent those colours seen in human vision. The HSI system has been selected and used successfully in three of the four machine vision systems that have been encountered. In any automated industrial application, the overall processing speed, however, is fundamental. The less pre-processing of initial input data, the better. In this light the RGB colour system would be the most suitable, as no coordinate system transformations are required:- the image is captured directly into RGB colour format using the RGB camera.

Although Hue would most certainly be able to quantify differences in colour between the different fruit classes, the success of the RGB colour model relative to the HSI model, was investigated. Typical cases of both fresh and dried fruit (good and reject) were compared within both the HSI and RGB colour model domains. The Hue, Saturation and Intensity values of dried peaches, pears and apricots, as well as fresh apples (red and yellow), purple plums and yellow pears were examined. The Red, Green and Blue values for these same fruits were observed.

Kay [9] found Hue and Intensity of the HSI colour system, to be successful in representing colour in all fresh apple types. Examination of various fruit types proved that the Hue and Intensity values adequately represented changes in both dried and fresh fruit colouring. In the RGB colour system, the change in the Blue value for fresh and dried fruit tones was either negligible (Blue values commonly approximately zero) or closely replicated the Green value or changes in Green value between fruit tones (colours). For this reason, only the Red and Green values of the RGB colour system were selected to represent the fruit colours. Reduction in the dimensionality of the feature space from three to two dimensions decreases the overall computation and processing time.

TABLE 1 consists of fruit examples that had values typical of that particular fruit type and fruit class. Two classes of fruit were examined : *good* and *bad*. Each set of fruit values (i.e. for *good* or *bad*), in the table, is the result of an average over 3 different fruits. For each fruit type, the fruit's average pixel values for the Red, Green, Hue and Intensity bands are shown (respectively R_{avg}, G_{avg}, H_{avg}, I_{avg}), as well as the respective average standard deviation values (R_σ, G_σ, H_σ, I_σ). The maximum value for any Red, Green or Blue value is 255. Although Hue values are typically measured by degrees, with 360° being the maximum Hue value, the Hue values in the table have all been normalised to the range $[0, 255]$; this makes visual comparison of the two models easier. Because Hue is commonly measured in a circular

fashion (0 to 360°), the shortest distance between two points may be measured clockwise or anticlockwise. TABLE 1 also lists the *Euclidean distance* (d_E) between the *good* and *bad* examples of each fruit type. For example, the *Euclidean distance* between the R_{avg} values of a good and a bad red apple in the (R,G) feature space is 115.

TABLE 1 : Average values and Euclidean distances(d_E) in R,G or H,I feature space

		Red and Green values				Hue and Intensity values		
FRUIT	STATS	GOOD	d_E	BAD	STATS	GOOD	d_E	BAD
RED	R_{avg}	170	115	55	H_{avg}	56	87	143
APPLE	R_σ	40	20	20	H_σ	53	14	67
	G_{avg}	25	22	3	I_{avg}	189	43	146
	G_σ	12	10	2	I_σ	18	17	35
PURPLE	R_{avg}	50	20	30	H_{avg}	181	29	152
PLUM	R_σ	20	10	10	H_σ	31	10	41
	G_{avg}	12	4	8	I_{avg}	145	4	141
	G_σ	3	1	2	I_σ	27	1	26
PEAR	R_{avg}	205	81	124	H_{avg}	18	2	20
(DRIED)	R_σ	10	20	30	H_σ	11	18	29
	G_{avg}	123	78	45	I_{avg}	211	11	200
	G_σ	5	10	15	I_σ	15	8	23
APRICOT	R_{avg}	180	80	100	H_{avg}	50	65	240
(DRIED)	R_σ	5	45	50	H_σ	16	19	35
	G_{avg}	60	30	30	I_{avg}	180	80	100
	G_σ	5	7	12	I_σ	4	46	50
PEACH	R_{avg}	170	50	120	H_{avg}	15	25	245
(DRIED)	R_σ	22	16	38	H_σ	5	3	8
	G_{avg}	60	24	36	I_{avg}	170	40	130
	G_σ	36	14	50	I_σ	12	28	40

If the *Euclidean distance* d_E between the *good* and *bad* fruit of a certain type is large, and the relevant standard deviations are small, then that feature space has differentiated the two fruit classes well. The two colour models can be compared therefore, by investigating which model achieves large *Euclidean distance* values, but small standard deviation values for the most fruit types. From the results in TABLE 1, it is clear that the R,G features would describe colour changes in fruit as adequately as would the H, I features: d_E values are comparable but in the case of the R_{avg} feature, mostly larger than any d_E distance in the corresponding H,I space; the standard deviation values are reasonably small in both feature spaces. Although the four

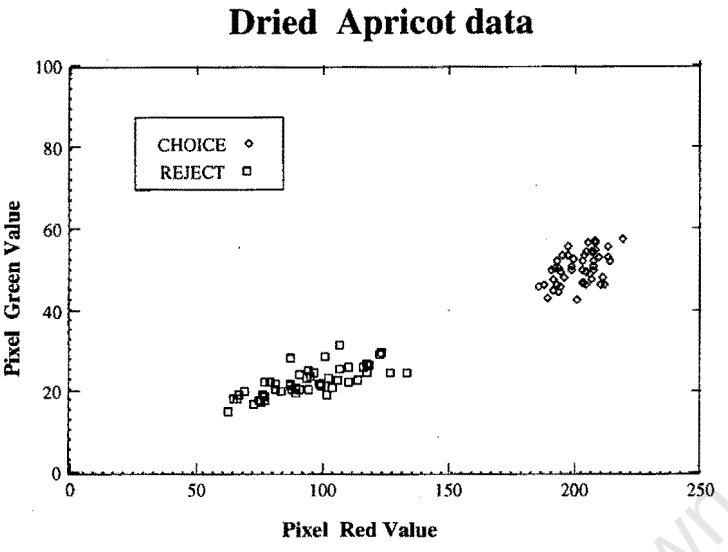

Figure 3.7: Scatterplot of the Red and Green band pixel values of a *good* and a *bad* dried apricot

features of R_{avg}, R_σ, G_{avg}, G_σ would describe the dried fruit colour changes well, a smaller dimensioned feature space was desirable. A 2-dimensional feature space of (R_{avg},G_{avg}) was investigated. This means that each dried fruit would be characterised by two values : R_{avg} would be the arithmetic average of all of the fruits Red band pixel values and G_{avg} would be the arithmetic average of all of the fruits Green band pixel values. R_{avg} and G_{avg} will be referred to as simply R and G hereafter.

Scatterplots in an (R,G) colour coordinate system for *good* and *bad* fruit types were examined to ensure reasonable class separabilities between (R,G) feature vectors in this feature space; an example is shown in Figure 3.7. The good spatial separation between the classes of *good* and *bad* dried fruit types on the scatterplots supported the decision to adopt an (R, G) feature vector to characterise the colour information in dried fruit. The 2-dimensional vector representation (R, G) was used in the final colour classification scheme. These R, G values were not normalised. Because the test images of the fruit were taken under strictly constant lighting conditions over a short period of time, little fluctuation in the lighting occurred; however, because the lighting may fluctuate over longer periods of time, it may be better to normalise the system data if the system was implemented.

In this chapter, a suitable colour system was investigated that would efficiently and accurately define the colour grades of dried fruit; various perceptual and quantitative colour models were investigated. Using available literature, an analysis of the colour models used in current fruit sorting machines, and a final investigation of RGB and HSI colour systems, it was concluded

that the (R,G) features of the (R,G,B) colour model adequately represented the range of colours within different dried fruit types and grades.

The next chapter deals with two of the segmentation problems that must be solved for efficient automatic dried fruit grading, in this machine vision system. They are (1) the isolation of the fruit from the image background, and (2) the isolation of Black blemishes on dried pears.

Section 3.5: Selection of a Colour Coordinate System

Chapter 4

Segmentation : Thresholding

4.1 Introduction to Segmentation

In order to examine specific regions of an image, it can be beneficial to segment the image into those regions. Segmentation is a process of classifying the set of given image pixels into subsets according to their distinguishing characteristics [51]. Rosenfeld and Davis [2] divide segmentation techniques into two basic classes of models: *statistical* models that describe the pixel population in an image or region, and *spatial* models that describe the decomposition of an image into regions. *Statistical* segmentation models describe the pixel population and its spatial distribution. *Spatial* segmentation methods include spectral and spatial classification [20], edge detection and edge-based methods [36], [1], [25], [49], [48], relaxation [15], thresholding [2], [48], morphology [8], and pattern classification and clustering [48]. Properties of the specific regions, or their relationship to other objects within the image, can be used to identify or highlight that region. The choice of segmentation method would include consideration of the *region description*. These are described by Rosenfeld [2] as :

1. **Interregion Interaction** : how the different regions or objects interact e.g.overlap
2. **Interregion Transitions** : whether object boundaries are blurred, sharp etc.
3. **Region size and shape** : what the shape and size of the objects are

There is no theory of image segmentation [53]. The methods of segmenting a region from an image are dependent entirely on the image characteristics, the available relationships between objects or areas within the image, and ultimately, the goals of the investigation. The final subsets will often be a result of a number of classifying procedures rather than one such classification. This is illustrated by Jain *et al* [24] who use a number of spatial segmentation

techniques to isolate muscle cell groups before classification; Jain *et al* [4] use both colour thresholding and texture analysis to isolate address labels on mail matter and Sun *et al* [58] use colour thresholding, iterative relaxation methods, filters and Bayes classification to separate cells for liver tissue classification.

In this system, segmentation is required for three tasks :

1. to isolate the dried fruit (object) from the image background, so that the dried fruit can be processed according to its own unique characteristics
2. to isolate the core area of a peach from the remainder of the peach flesh, so that the colour of the peach flesh can be more accurately analysed and classified
3. to accurately isolate any black blemishes present on dried pears.

Although all of the listed segmentation techniques were investigated, the two methods that were found to be the most successful with respect to the three tasks specified were: for problem (1) and (3), thresholding using histograms of the RGB colour bands, and for problem (2), morphology. In problem (1) the fruit (object) must be isolated from the background region. Grey level thresholding produces very good results in situations where the illumination gradient in the original raster image is regular and not too large [19]. As long as there is little intensity gradient in the colour of the background, histogram thresholding will be successful.

Using the colour feature constructively by ensuring the image background colour to be as different as possible from the general colours of dried fruit, would enhance any thresholding technique. The colour White was found to have high values of Red, Green, and Blue, and White is markedly different from the colours of dried fruit. (This is discussed further in the Appendix.) Jain *et al* [4], in fact, used this knowledge (of high R,G,B values in White) to successfully threshold white backed address labels from magazine covers. Tests which varied the local illumination level of the fruit, ensured that a fair range of Whites would still ensure successful thresholding using image colour histograms. The Blue band values of the RGB colour system showed particular success in this problem. The Blue band values of all dried fruit tested were very low, not much greater than zero. The Blue value of White, on the other hand, is very high; the Blue band therefore, was found to be the most ideal candidate for thresholding dried fruit from the background using histograms.

Problem (3) was a similar task to problem (1): Black blemishes were to be isolated from light golden yellow coloured pear flesh. Thresholding of colour histograms was again found to be an ideal solution. This particular problem, however, required the thresholding of both the Red and Green band values of the original RGB colour image. Section 4.5 covers the details of the segmentation of Black blemishes from dried pears.

The basic theory of Thresholding is described in Section 4.2. Five Automatic thresholding techniques were investigated, and are detailed in Section 4.3. Each of the described methods use the histogram of the image to select a threshold, and return a final optimum threshold value, T.

4.2 Thresholding

Image Thresholding is the process of segmenting an image into non-overlapping regions by isolating the image pixel values into different classes.

Consider a simple case : a binary output image $\mathbf{v_{out}}$ (containing pixel values of only 0 and 1) will be created if a threshold, T is applied to each pixel $\mathbf{v_{in}(i,j)}$ such that

$$\mathbf{v_{out}(i,j)} = \begin{pmatrix} 0 \text{ if } \mathbf{v_{in}(i,j)} < T \\ 1 \text{ otherwise} \end{pmatrix}$$

Extending this to many different classes with threshold values t_1, t_2, t_3 we have :

$$\mathbf{v_{out}(i,j)} = \begin{pmatrix} 0 & \text{if} & 0 \leq \mathbf{v_{in}(i,j)} < t_1 \\ 1 & \text{if} & t_1 \leq \mathbf{v_{in}(i,j)} < t_2 \\ 2 & \text{if} & t_2 \leq \mathbf{v_{in}(i,j)} < t_3 \end{pmatrix}$$

The threshold value can be selected manually or automatically.

4.2.1 Grey-level Thresholding

This technique of segmenting images is based on the pixel's grey level values. A colour image can be mapped onto a range of graded grey values using a lookup table [1](LUT). Each pixel in the image then has a designated grey value.

4.2.2 Grey-Level Histograms

We would sometimes like to have an idea or a measure of the distribution of the pixel values in an image. Histograms serve this purpose well. As highlighted by Liu [15] histogram-based segmentation has the advantage that it does not need *a priori* information of the image.

[1]The use of a lookup table (LUT) requires an entry for each possible input pixel value e.g. 256. The image's pixel values, once stored in a file can be mapped to another set of values using a Lookup table. LUT's ensure that the original file's pixel values can remain unchanged

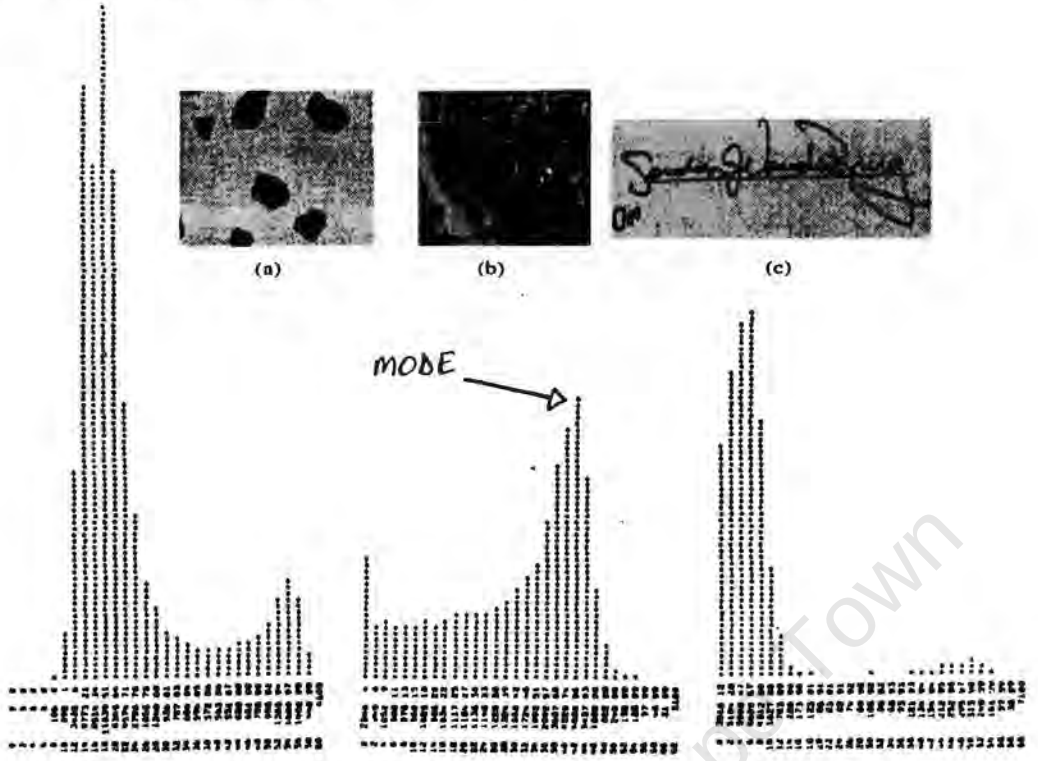

Figure 4.1: Grey-level histograms for the pictures (a)-(c)

A histogram of a grey level image is a measure of the pixel grey level distributions within the image and is derived by recording the number of pixels at each particular grey level. The domain of the histogram is the set of possible grey level values to which a pixel may be mapped. Three examples are illustrated in Figure 4.1. The image as well as the corresponding histogram is given for each example. Each histogram has three columns of data below it. The first column (0,2,4,..) represents the grey levels of the image where a value of 0 is pure white; the second column shows the number of pixels which have those grey level values, and the third column shows the percentile. The use of histograms is not limited to grey-scale values. For example the Red pixel values from an RGB image may be used to observe the distribution of Red values within an image.

The **Histogram Bin Size**, or rather bin width, is indicative of the number of consecutive grey levels included in that bin. For example, for a bin width of four, each bin contains information from pixels of four consecutive grey levels. Horn [6] discusses the trade-off between the histogram bin size and resolution. If the bins are too wide, the histogram resolution is low; if the bins are too narrow with bins containing only a few entries, the histogram will be irregular, making it difficult to ascertain a clear pattern of pixel frequency distribution.

The **mode of a histogram** is the value of the observation that occurs most frequently. A mode represents a peak value in a frequency distribution, and is a visible peak on the histogram, as indicated in Figure 4.1. A histogram may have more than one mode. Thresholds

may be defined by selecting the values that occur in the valleys of the histogram between the modes.

The size of the object of interest in the image, relevant to other objects or areas in the image is important. If, for example, the area of the object of interest within the image is small relative to the area of the background, then the peaks in the histogram will have vastly different dimensions. The respective contribution of these different areas to the histogram statistics are proportional to their size in pixels. The smaller peak, in this case, could become submerged in the tail of the larger peak and be rendered insignificant. This must be considered when selecting a suitable image segmentation procedure.

4.3 Automatic Thresholding Techniques

The five thresholding methods that are detailed in this section are the :

- Entropy Method [46]
- Moment-Preserving Method [52]
- Minimum-Error Method [14]
- Maximum Correlation Method [3]
- Discriminant Analysis Method [32]

In the following discussion of automatic thresholding techniques, an image f will contain N pixels and $(n+1)$ grey levels. (x, y) will be a pixel position in image f, and g will be that particular pixel's grey level. Each of the thresholding methods will calculate an optimum threshold value T. Each pixel value in image f will be above, below or equal to T. Image f will be segmented by T into two classes of pixels, those with values less than or equal to T, and those with pixel values above T.

4.3.1 The Entropy Method

Pun [46] proposed two iterative algorithms for threshold selection which developed a criterion function based on the entropy of the histogram. Kapur, Sahoo and Wong [18] developed a new thresholding algorithm, based on the same principles as Puns' method.

Using this method two probability distributions are defined for a selected threshold value. One probability distribution is for the grey levels which lie above the threshold value, the other for those grey levels which are less than or equal to the threshold value. The entropy

associated with each of the two distributions is calculated. The criterion function is the sum of these two entropies. The threshold value which corresponds to the maximum value, **T** of the criterion function, is the optimum one.

Let the grey levels of an image f be $g = 0, 1, .., n$. The grey level histogram can be viewed as a probability distribution where $f_0, f_1, ..f_n$ are the observed grey level frequencies in image f. The grey level probability distributions of image f are $p_0, p_1, .., p_n$ and are defined as :

$$p_g = \frac{f_g}{N}; \qquad g = 0, 1, .., n \tag{4.1}$$

$$N = \sum_{g=0}^{n} f_g; \qquad g = 0, 1, .., n \tag{4.2}$$

where N is the total number of pixels in the image, and there are $(n + 1)$ grey levels.
From the grey level probability distribution two probability distributions A and B, are derived; one for discrete values 0 to T, the other for values from $(T + 1)$ to n.

$$A : \frac{p_0}{P_T}, \frac{p_1}{P_T}, .., \frac{p_T}{P_T}$$

$$B : \frac{p_{T+1}}{1 - P_T}, \frac{p_{T+2}}{1 - P_T}, .., \frac{p_n}{1 - P_T}$$

where P_T and $(1 - P_T)$ are the class probabilities given by :

$$\begin{aligned} P_T &= \sum_{g=0}^{T} (p_g)\ (1 - P_T) \\ &= \sum_{g=T+1}^{n} p_g \end{aligned} \tag{4.3}$$

The entropy of the histogram is defined as :

$$H_n = -\sum_{g=0}^{n} p_g\ ln(p_g) \tag{4.4}$$

where p_g is the probability distribution of grey level g.

The entropy associated with distributions A and B are as follows :

$$\begin{aligned} H(A) &= -\sum_{g=0}^{T} \frac{p_g}{P_T} \; ln(\frac{p_g}{P_T}) \\ &= lnP_T + \frac{H_T}{P_T} \end{aligned} \tag{4.5}$$

$$\begin{aligned} H(B) &= -\sum_{g=T+1}^{n} \frac{p_g}{1-P_T} \; ln(\frac{p_g}{1-P_T}) \\ &= ln(1-P_T) + \frac{H_n - H_T}{1-P_T} \end{aligned} \tag{4.6}$$

where H_T is derived from Equation 4.4.

$$H_T = -\sum_{g=0}^{T} p_g \; ln(p_g) \tag{4.7}$$

The criterion function $\varphi(T)$ is defined as $H(A) + H(B)$:

$$\varphi(T) = ln(P_T) + \frac{H_T}{P_T} + ln(1-P_T) + \frac{(H_n - H_T)}{(1-P_T)} \tag{4.8}$$

The value T which maximises $\varphi(T)$ is selected as the optimum threshold.

4.3.2 The Moment Preserving Method

The Moment Preserving Method chooses the threshold in a way that preserves the first three moments of the original image, in the final thresholded bilevel image. Unlike the Entropy method, the threshold is obtained without the need for iteration.

Tsai [52] defines the i^{th} moment, m_i, of image f as :

$$m_i = \frac{1}{n} \sum_x \sum_y g^i; \qquad i = 1, 2, 3, .. \tag{4.9}$$

where g = pixels grey level at (x, y).

Moments can also be computed from the histogram of image f :

$$m_i = \sum_{g=0}^{n} p_g \, g^i$$

where $p_g = \frac{f_g}{N}$ given by equation 4.1.

Let the bilevel image, f_b consist of pixels with grey levels z_0 or z_1, where $z_0 < z_1$. Let P_0 and P_1, where $P_0 + P_1 = 1$, denote the fractions of pixels below or above the threshold value respectively. The first three moments of f_b are :

$$m_i' = \sum_{k=0}^{1} P_k z_k^i; \qquad i = 1, 2, 3 \tag{4.10}$$

Preserving the first three moments in f_b requires the following equalities :

$$m_i' = m_i; \qquad i = 1, 2, 3 \tag{4.11}$$

Because $P_0 + P_1 = 1$, the following four equalities can be derived :

$$\begin{aligned} P_0 z_0^0 + P_1 z_1^0 &= m_0 = 1 \\ P_0 z_0^1 + P_1 z_1^1 &= m_1 \\ P_0 z_0^2 + P_1 z_1^2 &= m_2 \\ P_0 z_0^3 + P_1 z_1^3 &= m_3 \end{aligned} \tag{4.12}$$

To find the optimum threshold T, the above four equations must be solved to find P_0 and P_1. T is then chosen to satisfy :

$$P_0 = \sum_{g=0}^{T} P_g \tag{4.13}$$

that is, the P_0-tile of the histogram. In the case where no discrete grey level exists which is exactly the P_0-tile of the histogram, the grey level closest to the P_0-tile is selected to be the threshold T.

The equations required to solve for P_0 and P_1 in 4.12 are detailed in [52].

4.3.3 The Minimum Error Method

In this method of Kittler and Illingworth [14], the grey level histogram is treated as a probability distribution. This method assumes that the grey level histogram is roughly bimodal, and that the two peaks are approximately Gaussian in shape. The distribution can therefore be approximated by the sum of two true Gaussian distributions. The threshold is used to divide the original distribution into two parts, each of which is then modelled by a Gaussian distribution. The threshold which produces the best fit, and consequently the least error, is chosen as the optimum threshold.

The two peaks have means of μ_0 and μ_1, standard deviations of σ_0 and σ_1 and *a priori* probabilities P_0 and P_1 respectively. Ideally, the grey level probabilities are given by :

$$p_g = \sum_{k=0}^{1} P_k \, p_{gk} \tag{4.14}$$

$$\text{where} \quad p_{gk} = \frac{1}{\sqrt{2\pi}\sigma_k} \exp\left[-\frac{(g-\mu_k)^2}{2\sigma_k^2}\right] \qquad k = 0, 1 \tag{4.15}$$

The 'minimum error threshold' is that grey level T for which the grey levels, $g = 0, 1, .., n$ satisfy

$$\begin{aligned} P_0 \, p_{g0} &> P_1 \, p_{g1} \qquad if \; g \leq T \\ P_0 \, p_{g0} &< P_1 \, p_{g1} \qquad if \; g > T \end{aligned} \tag{4.16}$$

Before T can be determined, the parameters μ_0, μ_1, σ_0, σ_1, P_0 and P_1 must be known. Using the technique of Kittler and Illingworth [14] these parameters are estimated from the grey level histogram as follows :

$$\mu_0(T) = \sum_{g=0}^{T} p_g \frac{g}{P_T} \tag{4.17}$$

$$\mu_1(T) = \sum_{g=T+1}^{n} p_g \frac{g}{1-P_T} \tag{4.18}$$

$$\sigma_0(T) = \sum_{g=0}^{T} (g-\mu_0(T))^2 \frac{p_g}{P_T} \tag{4.19}$$

$$\sigma_1(T) = \sum_{g=T+1}^{n} (g-\mu_1(T))^2 \frac{p_g}{1-P_T} \tag{4.20}$$

where P_T and $(1-P_T)$ are the class probabilities given by Equation 4.3, p_g are the grey level probabilities from Equation 4.1 and T is a candidate threshold.

The Gaussian distributions of p_{gk} now become :

$$p_{gk}(T) = \frac{1}{\sqrt{2\pi}\sigma_k(T)} \exp\left[-\frac{(g-\mu_k(T))^2}{2\sigma_k(T)^2}\right]; \qquad k = 0, 1 \tag{4.21}$$

The conditional probability, $e_g(T)$ of grey level g being replaced in the image by a correct binary value is :

$$e_g(T) = P_k(T)\frac{p_{gk}(T)}{p_g}; \qquad g = 0, 1, ..n \tag{4.22}$$

$$\text{where } k = \begin{cases} 0 & g \leq t \\ 1 & g > t \end{cases}$$

Since p_g is independent of both k and T, the authors ignore it in deducing a measure of classification error, $\varepsilon_g(T)$

$$\varepsilon_g(T) = \left[\frac{g-\mu_k(T)}{\sigma_k(T)}\right]^2 + 2\ ln\ \sigma_k(T) - 2\ ln\ P_k(T) \tag{4.23}$$

$$\text{where } k = \begin{cases} 0 & g \leq t \\ 1 & g > t \end{cases}$$

The average classification error for the whole image can be characterised by the criterion function :

$$J(T) = \sum_{g=0}^{n} P_g\ \varepsilon_g(T) \tag{4.24}$$

Substitution into $J(T)$ yields :

$$\begin{aligned} J(T) = 1 + 2[&P_T\ ln\ \sigma_0(T) + (1-P_T)\ ln\ \sigma_1(T)] \\ &-2P_T\ ln\ P_0(T) + (1-P_T)\ ln\ (1-P_T)] \end{aligned} \tag{4.25}$$

As the threshold T is varied, the Gaussian distributions change. The better the fit between the data and the distributions, the smaller the classification error becomes. The value T which yields the lowest value for $J(T)$ is then the minimum error threshold.

4.3.4 The Maximum Correlation Method

This thresholding technique, presented by Brink [3], involves direct comparison of the original image and the bilevel result. A statistical measure of correlation between the original image and the bilevel image is evaluated for different thresholds. The correlation will vary with threshold value, and a maximum correlation will indicate the optimum threshold.

Consider two sets of data :

$$\begin{aligned} X &= x_1, x_2, ..., x_n \\ Y &= y_1, y_2, ..., y_n \end{aligned}$$

Let X represent the set of grey level values g, of the individual pixels in the original image. Let Y represent those in the bilevel image. Let $p(x_i)$ and $p(y_i)$ be the probabilities of each data value x_i and y_i and $p(x_i, y_i)$ the probability associated with their product.

The coefficient of correlation, ρ (ref Freund(79) pg3-16) for the two sets of data X and Y is defined as :

$$\rho_{XY} = E_{XY} - \frac{E_X E_Y}{\sqrt{V_X V_Y}} \tag{4.26}$$

where E_X and E_Y are the expected values, or rather the mathematical expectations of the sets of data X and Y, and E_{XY} is the expected value of their product. V_X and V_Y are the variances of the sets X and Y. The expected values of X and Y, their product, and the variances are given by :

$$E_X = \sum_{i=1}^{n} x_i \; p(x_i) \tag{4.27}$$

$$E_Y = \sum_{i=1}^{n} y_i \; p(y_i) \tag{4.28}$$

$$E_{XY} = \sum_{i=1}^{n} x_i \; y_i \; p(x_i, y_i) \tag{4.29}$$

$$V_X = E_{XX} - [E_X]^2 \tag{4.30}$$

$$V_Y = E_{YY} - [E_Y]^2 \tag{4.31}$$

$$\text{where} \quad E_{XX} = \sum_{i=1}^{n} (x_i)^2 \; p(x_i) \tag{4.32}$$

$$E_{YY} = \sum_{i=1}^{n} (y_i)^2 \; p(y_i) \tag{4.33}$$

The correlation coefficient, ρ_{XY} can take on values $(-1, 1)$ where $\rho_{XY} = 0$ indicates no correlation at all between X and Y, $\rho_{XY} = -1$ indicates an inverse correlation, and $\rho_{XY} = 1$ indicates an absolute correlation.

The correlation criteria can be applied to an image as follows :

If an image is made up of $(k \times \ell)$ pixels, then each pixel constitutes $1/N$ of that image where $N = (k \times \ell)$. The two levels of the bilevel image Y are represented by the below-threshold mean $\mu_0(T)$ and above-threshold mean $\mu_1(T)$ of the original image X as given by equation 4.17 and 4.18 in Section 4.3.3.

The coefficient of correlation between the two images X and Y for threshold T is :

$$\rho_{XY}(T) = E_{XY}(T) \; - \; \frac{E_X \; E_Y(T)}{\sqrt{V_X \; V_Y(T)}} \qquad 0 \leq T < n \tag{4.34}$$

$$\begin{aligned}
\text{where} \quad E_X &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} x_{ij} \; \frac{1}{N} \\
E_Y(T) &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} y_{ij} \; \frac{1}{N} \qquad 0 \leq T < n \\
E_{XY}(T) &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} x_{ij} \; y_{ij} \frac{1}{N} \qquad 0 \leq T < n \\
E_{XX} &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} (x_{ij})^2 \frac{1}{N} \\
E_{YY}(T) &= \sum_{i=1}^{k} \sum_{j=1}^{\ell} (y_{ij})^2 \frac{1}{N} \qquad 0 \leq T < n
\end{aligned}$$

with the variances given by :

$$\begin{aligned}
V_X &= E_{XX} - [E_X]^2 \\
V_Y(T) &= E_{YY}(T) - [E_Y(T)]^2 \qquad 0 \leq T < n
\end{aligned}$$

Only E_X and V_X are independent of the threshold T, as they are obtained directly from the original image X. The threshold which maximises the correlation function $\rho_{XY}(T)$, is found

by an iterative process and is selected as the optimum threshold.

4.3.5 The Discriminant Analysis Method

The method of thresholding discussed by Otsu [32] is discussed here.

This threshold operation divides the pixels of the original image into two classes; C_0 for the object and C_1 for the background with threshold value T. C_0 thus represents pixels with grey levels $(0, 1, .., T)$,and C_1 represents pixels with grey levels $(T+1, T+2, ..n)$.

The image histogram is normalised and regarded as a probability distribution :

$$p_g = \frac{f_g}{N}$$

where

$$\begin{cases} f_g = & \text{frequency of occurrence of grey level g} \\ N = & \text{number of pixels in the image} \end{cases}$$

The probabilities of class occurrence ω are :

$$\begin{aligned} P_r(C_0) &= \omega_0 = \sum_{g=0}^{T} p_g \\ P_r(C_1) &= \omega_1 = \sum_{g=T+1}^{n} p_g \end{aligned} \tag{4.35}$$

The mean levels μ for each class are :

$$\mu_0 = \sum_{g=0}^{T} g\ P_r(\frac{g}{C_0}) = \sum_{g=0}^{T} g\ \frac{p_g}{\omega_0} = \frac{\mu(T)}{\omega(T)}$$

$$\mu_1 = \sum_{g=T+1}^{n} g\ P_r(\frac{g}{C_1}) = \sum_{g=T+1}^{n} g\ \frac{p_g}{\omega_1} = \frac{\mu_{TOT} - \mu(T)}{1 - \omega(T)}$$

$$\text{where} \quad \begin{aligned} \mu(T) &= \sum_{g=0}^{T} g\ p_g \\ \omega(T) &= \sum_{g=0}^{T} p_r \end{aligned} \tag{4.36}$$

$\omega(T)$ and $\mu(T)$ are the *zeroth* and *first order* cumulative moments of the histogram up to grey level T and

$$\mu_{TOT} = \mu(n) = \sum_{g=0}^{n} g\ p_g$$

is the total mean grey level of the original image.

Otsu [32] states that for any value of T :

$$\omega_0 + \omega_1 = 1 \tag{4.37}$$

$$\omega_0\ \mu_0 + \omega_1\ \mu_1 = \mu_{TOT} \tag{4.38}$$

The class variances are given by :

$$\begin{aligned}(\sigma_0)^2 &= \sum_{g=0}^{T}(g-\mu_0)^2 P_r(\frac{g}{C_0}) \\ &= \sum_{g=0}^{T}(g-\mu_0)^2\ \frac{p_g}{\omega_0} \qquad (4.39)\\ (\sigma_1)^2 &= \sum_{g=T+1}^{n}(g-\mu_1)^2\ P_r(\frac{g}{C_1}) \\ &= \sum_{g=T+1}^{n}(g-\mu_1)^2\ \frac{p_g}{\omega_1} \qquad (4.40)\end{aligned}$$

Measures of class separability or *Discriminant Criterion Measures* are used to evaluate the optimum threshold T :

$$\lambda = \frac{\sigma_B^2}{\sigma_W^2}\ ;\quad \kappa = \frac{\sigma_{TOT}^2}{\sigma_W^2}\ ;\quad \eta = \frac{\sigma_B^2}{\sigma_{TOT}^2} \tag{4.41}$$

$$\begin{aligned}\text{where}\quad \sigma_B^2 &= \omega_0(\mu_0-\mu_{TOT})^2 + \omega_1(\mu_1-\mu_{TOT})^2 \\ &= \omega_0\omega_1\ (\mu_1-\mu_0)^2 \qquad (4.42)\\ \sigma_W^2 &= \omega_0(\sigma_0)^2 + \omega_1(\sigma_1)^2 \qquad (4.43)\\ \sigma_{TOT}^2 &= \sum_{g=0}^{n}(g-\mu_{TOT})^2\ p_g \qquad (4.44)\end{aligned}$$

$(\sigma_B)^2$, $(\sigma_W)^2$ and $(\sigma_{TOT})^2$ are respectively the between-class variance, the within-class vari-

ance and the total variance.

Otsu [32] defines the maximum threshold T, as that value which gives the best separation of classes in grey levels. The problem is thus one of finding a threshold value T which maximises one of the criterion measures λ, κ, or η.

It is assumed that the following relation always holds :

$$(\sigma_W)^2 + (\sigma_B)^2 = (\sigma_{TOT})^2 \tag{4.45}$$

which implies the discriminant criteria maximising of λ, κ, or η of Equation 4.41 are equivalent.

Unlike $(\sigma_B)^2$ which is based on first-order statistics(class means), $(\sigma_W)^2$ is based on second-order statistics. $(\sigma_{TOT})^2$ is independent of T. η is therefore selected as the simplest criterion measure with respect to T.

The optimal threshold, therefore, is selected using :

$$\eta(T) = \frac{\sigma_B(T)^2}{\sigma_{TOT}^2} \tag{4.46}$$

but as $(\sigma_{TOT})^2$ is independent of T, only $\sigma_B(T)^2$ need be considered. $\sigma_B(T)^2$ can be written :

$$\sigma_B(T)^2 = \left[\frac{\mu_{TOT}\,\omega_T - \mu_T^2}{\omega_T[1-\omega_T]}\right] \tag{4.47}$$

The optimal threshold, T, occurs when $\sigma_B(T)^2$ is a maximum.

4.4 Application : Segmentation of the Background from the Image

The five automatic thresholding techniques detailed in this chapter were examined in order to find the automatic thresholding technique which best thresholded the object (fruit) from the White coloured (see Appendix for discussion on the selection of the background colour) background. Each of the thresholding techniques makes use of the image histogram. This means that only one band or coordinate of a colour system can be thresholded. If, for example, a grey-scale version of a dried fruit image was most effective for thresholding, then the Intensity coordinate or band of the HSI colour system, would be a suitable choice. For dried fruit images upon a White background, the Blue band values (from the RGB colour coordinate system) were found to produce well separated modal peaks in the histogram. The

Section 4.4: Application : Segmentation of the Background from the Image

Blue band values were selected for the process of thresholding the fruit from the background of the image.

A 'best choice' threshold value was selected manually for a number of test images. Each of the five automatic thresholding techniques were then applied to these images, and the results compared to their manual 'best choice' threshold value. The thresholding method which produced the threshold value closest to the visual, or manual threshold value, was selected as the best method for that fruit type. The results are listed in TABLE 1. Discriminant Analysis (DA) was clearly the most successful method. In the cases where the Minimum Error (ME) method was favoured, the difference between the DA and ME threshold value was very small. In Figure 4.2 the success of the Discriminant Analysis method is seen; clearly each dried fruit has been removed very successfully from the background.

TABLE 1 : Automatic Threshold methods - Results

Dried Fruit type	Discriminant Analysis	Minimum Error	Maximum Correlation	Moment Preserving	Entropy
Choice Pear	√				
Apple	√				
Banana	√				
Choice Apricot	√				
Choice Peach		√			
Bad Pear		√			
Bad Apricot	√				
Green Peach	√				
TOTAL	6	2			

The success of each method can be predicted according to the type of image histogram; one would expect image data that produce reasonably clear bimodal histogram peaks to be well thresholded by, for example the Minimum Error method, which assumes a reasonably bimodal histogram. The current data of the Blue band values of a dried fruit upon a White background, fall into this category of clear bimodal peaks. An important consideration, in this case, is that the object (fruit) in the image is not too small relative to the expanse of the background.

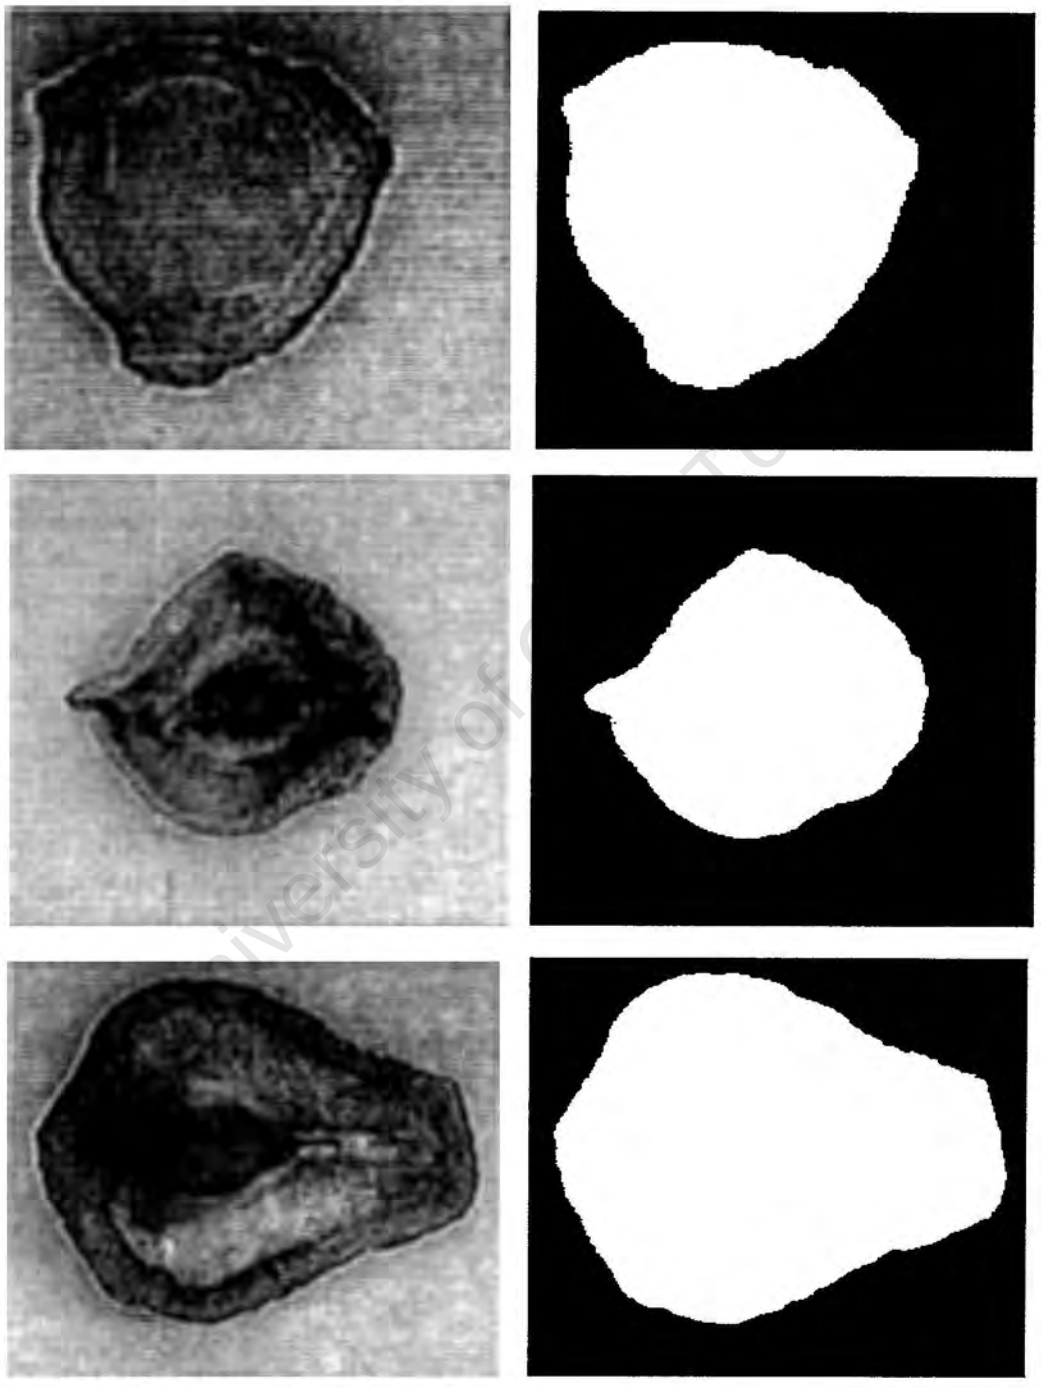

Figure 4.2: Thresholding of the background from various dried fruit

4.5 Application : Identifying and Isolating Blemishes in Dried Pears

Dried pears are graded into four classes by South African Dried Fruit (SADF). The 3rd and 4th classes consist of those dried fruit that have Black blemishes. The pears allocated to Class 4 have blemishes that are larger than those of the pears allocated to Class 3. The remainder of the pear flesh on these blemished pears, is normally identical in colour, to that of Class 1 pears. Dried pears that are allocated to Class 3 or 4, are allocated by a knowledge of blemish existence and blemish size alone.

Each of the three bands of the RGB image were investigated in order to find that band which achieved the best threshold results. The Red and Green colour bands, when individually thresholded and the resulting masks combined, produced the best thresholding of Black blemishes from dried pear flesh.

Again, visual examination of the best manual thresholding value was used to discover the most appropriate automatic thresholding technique for the task. The Entropy automatic thresholding technique was found to be the most successful in isolating all blemishes on a dried pear. This method does not assume a bimodal distribution, and produced thresholding results far superior to those of the other four methods.

Process of Blemish Isolation

Two histograms were used to establish a mask of all blemish areas on a pear. One was the histogram of the fruit pixels Red values, the other was the histogram of the fruit pixels Green values. Both histograms were automatically thresholded using the Entropy method. The two thresholding procedures produced two binary images i and j. The binary images i and j were then combined as follows to produce a binary image k :

$$k = i + j \qquad \text{where } + \text{ is a LOGICAL OR procedure}$$

This combination of the two binary images ensured that blemishes isolated in both the Red and Green bands were found. It was found that blemishes found by thresholding either the Green or Red band alone were not entirely successful; combining the results from the two resulted in excellent blemish isolation on dried pears. Both the grey-scale images of three blemished pears and the binary images resulting from the thresholding of both the Red and Green value histograms are shown in Figure 4.3. It should be noted that dried pear blemishes are distinctly Black in colour, rather than Brown as seen in other types of

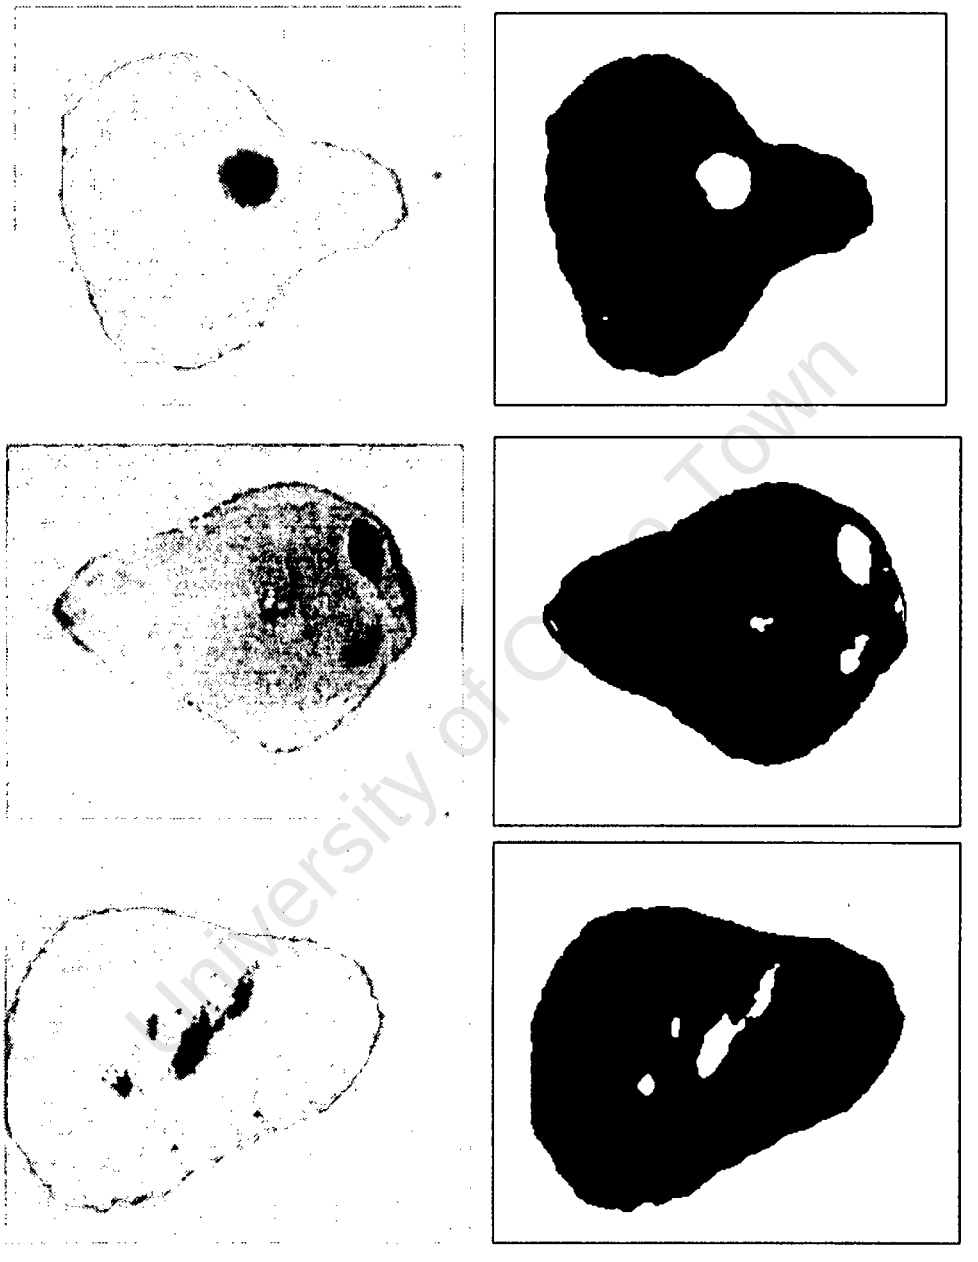

Figure 4.3: Thresholding of black blemishes from dried pears

dried fruit. Sometimes a Yellow blemish area, rather than a Black blemish is isolated by this thresholding procedure. In order to ignore these cases, the threshold value T selected by the Entropy procedure can be compared to the histogram peak or maximum value. Because the main peak, which describes the pear flesh, has a Red or Green value lower than the yellow blemish peak, one can ignore blemishes where the threshold value T is greater than the histogram maximum value.

In many cases, the thresholding procedure highlighted the dark border line of the fruit; this area was erroneously included with the blemishes in a binary mask. This problem was dealt with effectively by doing the following procedure :

Threshold the background of the image from the initial fruit image using the method described above. The border of the binary mask of the fruit image can then be tracked using the method investigated and implemented by Bartleet [44]. This boundary is then *dilated* by 3 pixels, and subtracted from the full colour image of the fruit. This ensures that the misleading darker border line of the fruit image is removed from the colour and blemish analysis processes.

Once the blemishes have been isolated they must be *labeled*, so that the pixel area of each blemish can be calculated. The existence of blemishes, as well as their pixel area (or diameter), is now available to be used in a dried fruit classifier to aid the colour classification scheme.

Two of the three segmentation tasks have been solved using automatic thresholding techniques, based upon the image histogram. Segmentation of the core area of a dried peach from the remaining flesh was not well facilitated by thresholding. Concentrating on the geometry of the problem (where the physical removal of the core or pip results in an indentation in the peach) resulted in an investigation and final solution in Morphology. Morphology and relevant Morphological operations are introduced in Chapter 5.

Chapter 5

Segmentation : Morphology

5.1 Introduction

In the proposed dried fruit sorting system, one of the primary tasks was to deal with the problem of grading dried peaches. Because the core or pip of a peach is small relative to the depth of the flesh of a peach (unlike an apricot where the pip is large relative to the depth of the flesh), its removal from the peach half leaves an indented area in the centre of the peach flesh. This indented area is dark coloured, and proved to be the same colour as a peach blemish. Various spatial segmentation techniques such as gradient, thresholding, and texture were investigated; none provided acceptable results. The geometry of the indentation was then investigated. Segmentation of the Edge-faces of the core area were considered: De Floriani [28] discusses the construction of a generalised edge-face graph (GEFG) to represent the topological information of an object. This involved indexing, connecting, and representing the topological information of an object. The boundary of the core area could be found using this method, but the result would be unnecessarily detailed and the entire process very time-consuming. A more time-efficient solution was required. (The exact topological nature of the peach would, however, be useful in finding whether a core still remained in the peach or not, a problem that does currently exist).

Mathematical Morphology can rigorously quantify many aspects of the geometrical structure of signals in a way that agrees with human intuition and perception [33]. For this reason a number of morphological operations were applied to this segmentation problem. The Top-Hat Transform, which is a combination of two of the primary operations of mathematical morphology, namely *erosion* and *dilation*, was found to be very successful in the segmentation of the core region from any given dried peach.

Section 5.2: Binary Mathematical Morphology

The following discussion is intended as an outline of the relevant concepts of morphological image processing as detailed by Dougherty [8], rather than a rigorous mathematical analysis. A more expansive and thorough discussion of Morphological operations and concepts can be found in [53]. The basic theories and concepts of Morphology, which support the Top-Hat Transform operation, are introduced in Section 5.2. For the sake of clarity, binary morphology was used to introduce the operations, although the solution required grey-scale morphology principles. An explanation of the Top-Hat transform is given in Section 5.4. The final process of segmenting the core area in an image from the peach using the Top-Hat transform, is described in Section 5.5.

5.2 Binary Mathematical Morphology

Mathematical Morphology refers to a branch of non-linear image processing and analysis that concentrates on the geometric structure within an image [8]. The scope of morphological methods is very wide; they include segmentation, edge detection, texture analysis, compression, curve filling and shape analysis.

The primary operation in Mathematical Morphology is to probe an image with a structuring element and quantify the manner in which it fits. Both the manner and extent of fit depend upon the size and shape of the structuring element [8]. There are two primary operations in Morphological image processing :- *erosion* and *dilation*. From erosion and dilation the morphological operations of *opening* and *closing* can be composed. It is these latter two operations that have a close connection to shape representation, decomposition and primitive extraction [38].

Only Euclidean binary images, which are a subset of n-dimensional Euclidean space, will be considered; in digital images, only a 2-dimensional or binary space is required.

5.2.1 Erosion

The nature of probing is to mark the translations of a structuring element and where it fits into the image. In binary mathematical morphology,the translation of a set A by a point x, is denoted by $(A + x)$ and is defined as :

$$A + x = \{a + x : a \in A\}$$

The *erosion* of set A by set B, written $A \ominus B$, is defined as :

$$A \ominus B = \{x : B + x \subset A\}$$

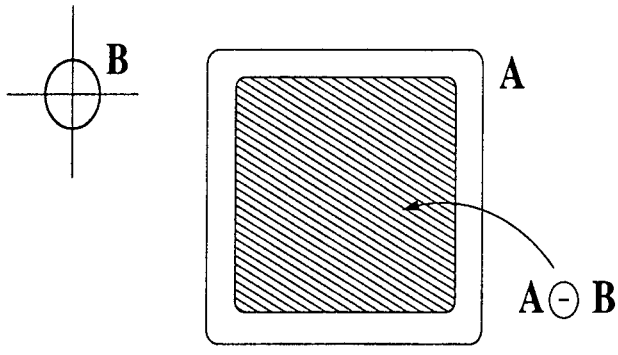

Figure 5.1: The process of Erosion where the structuring element B is a disk

where $\subset$ denotes subset relation, B is the structuring element and B is centred at the origin. A $\ominus$ B therefore consists of all points x for which the translation of B by x fits inside of A.

In Figure 5.1 the structuring element selected is a disk with the origin at the disk centre. Geometrically, the disk is 'rolled' around the inside of A and the positions of the origin (in this case the disk centre) are marked to produce the eroded image. Should the origin of the structuring element not lie within the structuring element, it may be that the eroded image will not lie within the input image.

Erosion can be represented by an intersection of image translates :

$$A \ominus B = \cap\{A - b : b \in B\}$$

In this formulation erosion is found by intersecting all translates of the input image by negatives of points in the structuring element [8].

5.2.2 Dilation

Dilation is the dual operation to *erosion* and is defined by set complementation of *erosion*. The *dilation* of set A, by the structuring element B, written $A \oplus B$, is defined as :

$$A \oplus B = [A^C \ominus (-B)]^C$$

where A^C denotes the complement of A.

To *dilate* A by B, B is rotated around the origin to obtain $(-B)$. A^C is then eroded by $(-B)$ and finally the complement of the erosion is taken. In Figure 5.2, the final result of dilation of set A by a disk B is an expansion of A. *Dilation* can also be represented by a translation

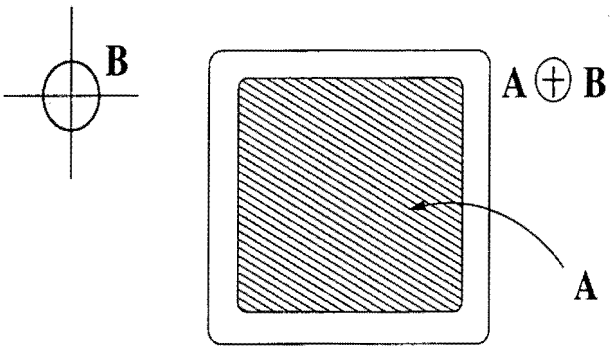

Figure 5.2: The process of Dilation where the structuring element B is a disk

of the input image by all points in the structuring element B, followed by taking the union thereof:

$$A \oplus B = \cup \{A + b : b \in B\} \quad .$$

This notation has historically been called **Minkowski** addition [8]. The difference in the effects of *erosion* and *dilation* operations upon an image can be seen in Figure 5.3 where the structuring element E is a 4-pixel square, with the lower left pixel situated at the origin and:

- The dilation has resulted in an expansion of the input image with small intrusions into the image being filled.
- Erosion has the expected 'shrinking' effect, with small extrusions in the image being eliminated. As the structuring element E is too thin to fit both the vertical piece at the top of the image, and the centre piece of the image, both have been eliminated in the erosion, leaving a disconnected image.

5.2.3 Opening

Both *opening* and *closing* operations are defined in terms of the *erosion* and *dilation* operations.

The *opening* of an image A by a structuring element B, written $A \odot B$, is defined as an *erosion* operation followed by a *dilation* operation.

$$A \odot B = (A \ominus B) \oplus B$$

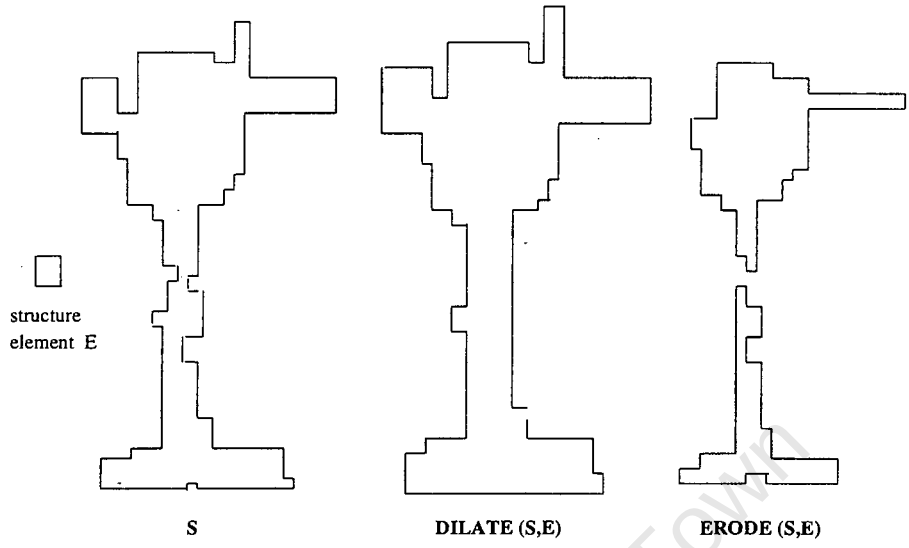

Figure 5.3: The effects of Digital Erosion and Dilation

An equivalent formulation would be :

$$A \odot B = \cup\{B + x : B + x \subset A\}$$

In this interpretation the opening results from a union of all translations of the structuring element that fit inside the input image.

5.2.4 Closing

Closing is the dual operation to *opening* - a *dilation* followed by an *erosion*. The closing of A by structuring element B, is written $A \bullet B$, and is defined :

$$A \bullet B = [A \oplus (-B)] \ominus (-B)$$

In Figure 5.4 the initial set S used in Figure 5.3 has the operations of open and close applied to it. As can be seen, the opening of the image has eliminated small extrusions in a much finer manner than the erosion operation. The opened image is a far better replica of the original image than is the eroded image). Analogously the closed image has a finer filling of image intrusions and is a better overall replica of the original image.

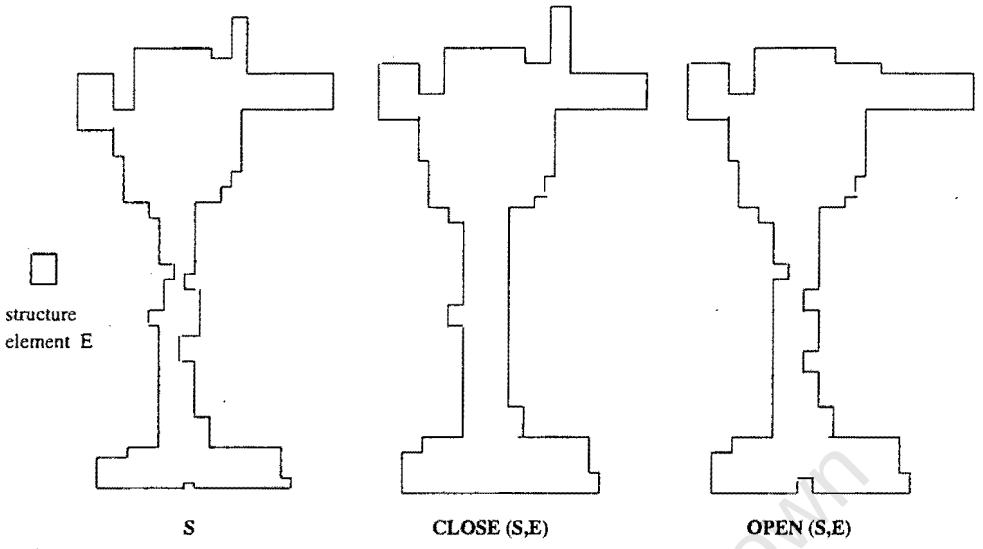

Figure 5.4: The effects of Digital Opening and Closing

5.3 Grey-Scale Morphology

In the late 1970's a transcription of all basic binary morphology operations was made to apply the same concepts to gray-scale images. Erosion and dilation can be performed on a grey-scale image and are equivalent to the computation of the minimum and maximum of the values in the local neighbourhood [33]. Grey-scale morphological processing is concerned with the topography of an images graph. Grey-scale morphology differs from binary morphology in that the morphological operators act on real valued functions defined on n-dimensional Euclidean space, rather than binary functions [8].

As with binary morphology, we need to define a grey-scale translation $(f_x + y)$:

$$(f_x + y)(z) = f(z - x) + y$$

This translation includes the horizontal shift of a signal f to the right by x :

$$f_x(z) = f(z - x)$$

and a vertical offset by y :

$$(f + y)(z) = f(z) + y$$

The counterparts of *intersection* and *union* in binary morphology, are *minimum* and *maximum* in grey-scale morphology. Given two signals f and g, we define the minimum of the two in

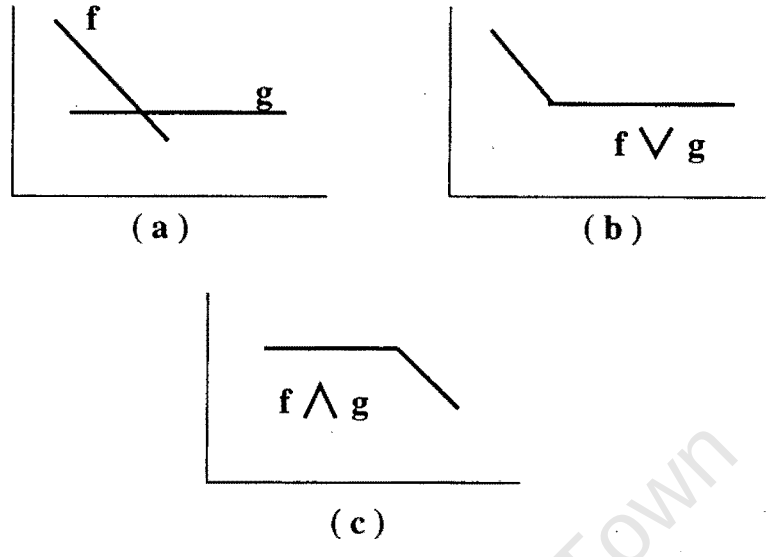

Figure 5.5: Maximum and Minimum operations on signals f and g

the following manner : if x lies in the intersection of the domains, $D[f] \cap D[g]$, then

$$(f \wedge g)(x) = min\{f(x), g(x)\}$$

otherwise, x does not lie in the domain of $f \wedge g$. Similarly the maximum $f \vee g$, for x lying in $D[f] \cap D[g]$, is defined pointwise by :

$$(f \vee g)(x) = max\{f(x), g(x)\}$$

The operations of Open, Close, Erode and Dilate are conceptually the same in grey-scale morphology as that of binary morphology.

5.4 The Top-hat Transform

In grey scale images, if you subtract an *opened* image from the original image, you will be left with features such as peaks and points of high curvature. In grey-scale morphology this would be useful in pinpointing certain markers in an image. These markers could include :

- small pixel clusters that are dark on a relatively lighter background, or
- edges in a reasonably noise free image.

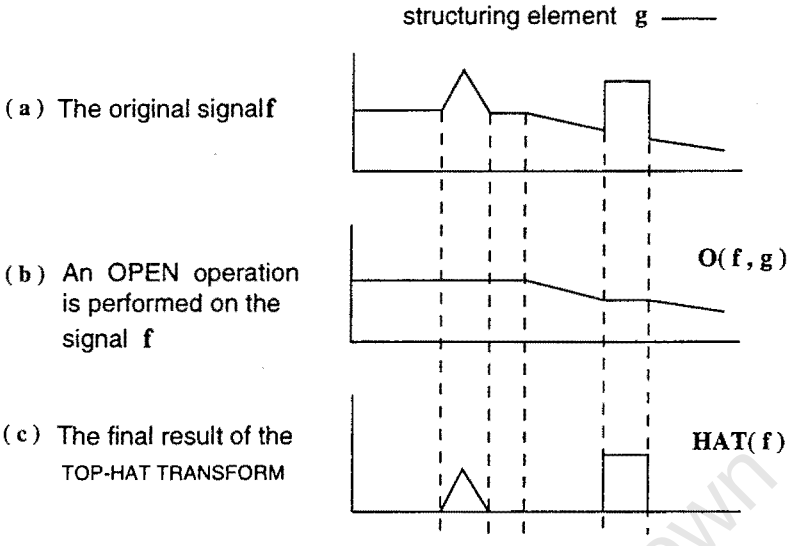

Figure 5.6: The Top-Hat Transform process applied to a signal f

The operation of subtracting an *opened* image from the original image has been labelled the **Top-Hat Transform** and is defined by :

$$\mathrm{HAT}(f) = f - (f \odot g)$$

where g is the chosen structuring element. Opening is an *antiextensive* [1] operation, therefore the opened image lies beneath the original image and $HAT(f)$ is always nonnegative.

A simple example demonstrates the Top-Hat transform operation in Figure 5.6. The structuring element g is flat with a length that is longer than the bases of the peaks that jut above in the original signal. The *open* operation has smoothed the baseline of the original signal. When the *opened* signal is deducted from the original signal, only the peaks in the original signal remain. The Top-Hat transform has caused the peaks to be highlighted and removed from the baseline of the signal.

Opening and closing operations offer an intuitively simple and mathematically formal way for peak or valley detection [33]. Maragos [33] illustrated the success of the Top-Hat transform in extracting features such as eyes and mouth from images of human faces. The indented region of a peach core can be compared to a valley; Maragos' investigation therefore, confirms the suitability of the Top-Hat transform for this problem.

[1] An operator F is said to be *antiextensive* if $F(A)$ is always a subset of A

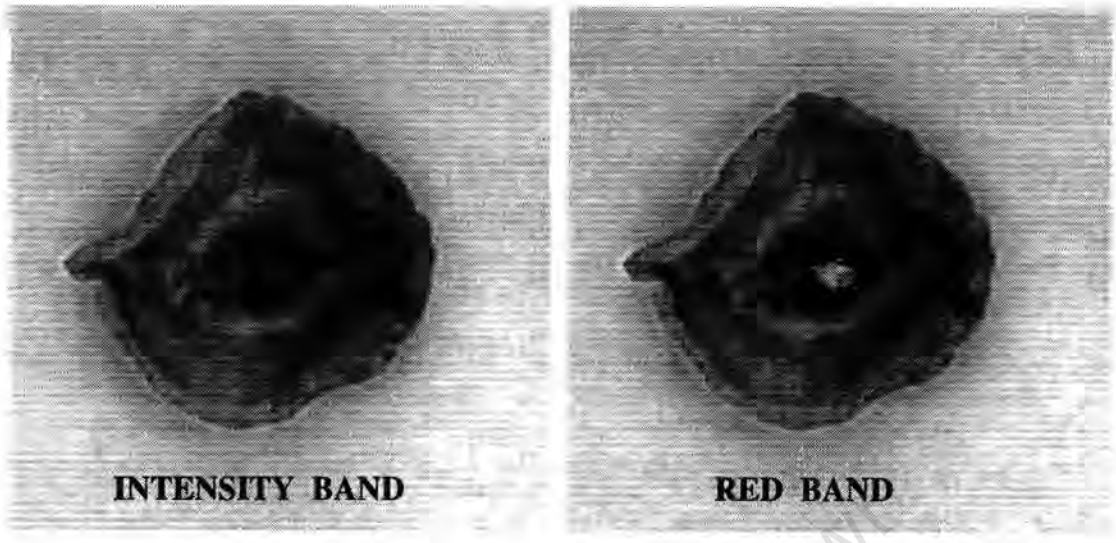

Figure 5.7: The Intensity and Red colour bands for a dried peach image

5.5 Application of the Top-Hat Transform to remove the core area from a dried peach

5.5.1 Introduction

The core of a fruit is removed before the fruit is sulphur dried. In the case of a peach, the removal of the core results in a 'core area' which is a much darker red, brown than the remainder of the peach flesh. This dark colouring causes problems with most dried fruit colour image processing systems, as the 'core area' is similar in colour to a blemish. In any successful dried fruit colour classification procedure, the dried peach core area will have to be examined independently from the remainder of the peach flesh.

5.5.2 The Top-Hat Transform operation

In Figure 5.7 images of the front of a dried peach in both the Intensity (I) band of the HSI colour coordinate system, and the Red (R) band of the RGB colour coordinate system are shown. In the Intensity band example of Figure 5.7, every pixel value in this image is the corresponding grey-scale value (commonly $I = \frac{1}{3}R + \frac{1}{3}G + \frac{1}{3}B$) of the original full 24 bit colour pixel value.

The Red band pixel values of the peach closely follow the rise and fall in the corresponding Intensity (I) pixel values. High Intensity values corresponding to lighter colours are reflected in high Red values; similarly low Intensity values corresponding to darker colours are reflected

Section 5.5: Application of the Top-Hat Transform to remove the core area from a dried peach

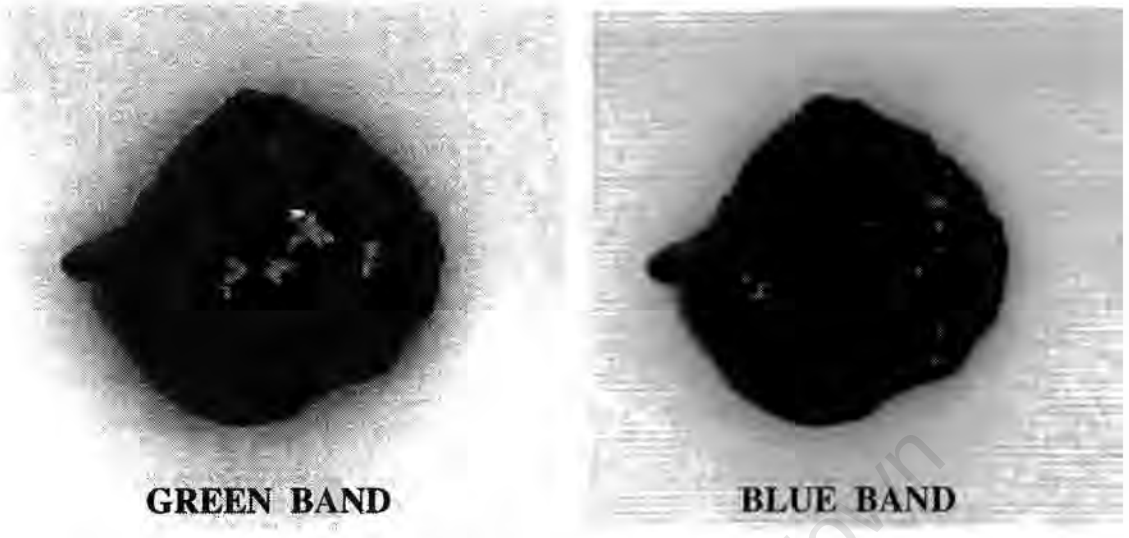

Figure 5.8: The Green and Blue bands of the RGB image of a dried peach

in low Red values. The close correlation between Red band and Intensity band image values is attributable to the consistently low Blue (approximately zero) and Green band values in dried peaches. The Red band values closely reflect the colour changes in a dried peach, and consequently accentuate the geometrical 'indentation' or valley of the core region. The Red band values of the RGB image were selected as the set to which the Top-Hat transform operator would be applied. Because it was not necessary to convert the RGB image to grey scale format before applying the transform, the overall segmentation procedure is much faster than it would otherwise have been. For comparison purposes, the green band image and blue band image of the same peach are shown in Figure 5.8.

In Figure 5.9 only the region of the image which enclosed the core area was selected. This region is displayed in a 3-dimensional colour contour plot of the Red band values. As noticed earlier, the Red values decrease in darker colours, notably within the core area; this results in a reasonably flat 'blanket' of similar or close Red values surrounding the core, and a very noticeable hollow indentation (reflecting the low Red values) depicting the dark coloured core. Similar results were seen in all images of dried peaches that were examined.

An operator was required to successfully isolate only the dark core area of the fruit: the Top-Hat transform was successful. However, if the Top-Hat transform was applied to the entire surface of the peach, it highlighted select blemishes as well as the required core area. To prevent this, the central region of the fruit where the core was always found, needed to be isolated. In each case the centre-of-mass of the peach was found, and only a square region about that centre-of-mass was processed.

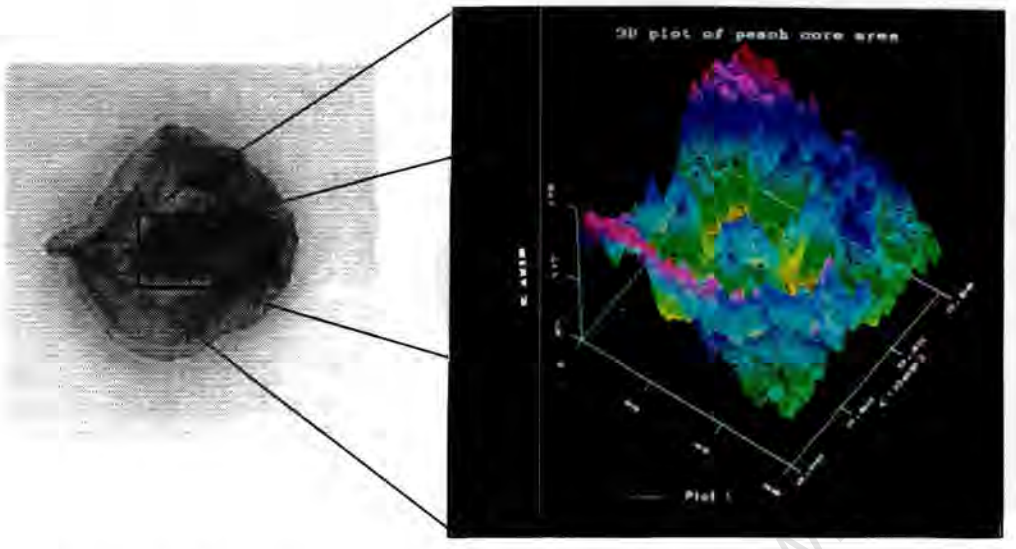

Figure 5.9: A 3-dim contour plot of the Red-band values in a peach core area

always successfully highlight any core area. In order to highlight the 'valley' of the core, the structure element had to have a diameter larger than the diameter of the core (valley) it was attempting to isolate. Three structuring elements were tested; a square, a sphere and a flat disk. A disk of suitable radius was the most successful; it isolated the core area thoroughly in the shortest time. The radius of the disk was dependent on the size of the image core pixel area that it was applied to : the original images that were captured consisted of four peaches, where each peach occupied a 240×240 pixel area. The centre-of-mass of each peach was found using an algorithm by Levine [31]. A pixel area or rather a window area of 120×120 about the centre-of-mass was then processed using the Top-Hat transform. The process remained completely successful but was much faster if the 120×120 image was shrunk by a magnitude of 8 before applying the Top-Hat transform. A disk shaped structure-element of radius of 5 pixels was found to be the most successful on these 15×15 pixel areas. The result of the Top-Hat transform on three different dried peaches, each from a different grade, is shown in Figure 5.10.

Once the core 'valley' had been identified by the transform, a grey level threshold had to be selected to generate a mask for the core area. A preset global threshold value could be used, but the results would then be very sensitive to any fluctuation in the lighting conditions. The Discriminant Analysis automatic thresholding algorithm described in Chapter 4 was found to be very successful for this purpose.

It is necessary to perform the core removal procedure only upon images of the 'front' of a dried peach. Simple statistical tests (see Appendix) of the fruit image indicate whether or not the image is a representation of the 'front' or 'back' of the peach. This ensures that no

Section 5.5: Application of the Top-Hat Transform to remove the core area from a dried peach

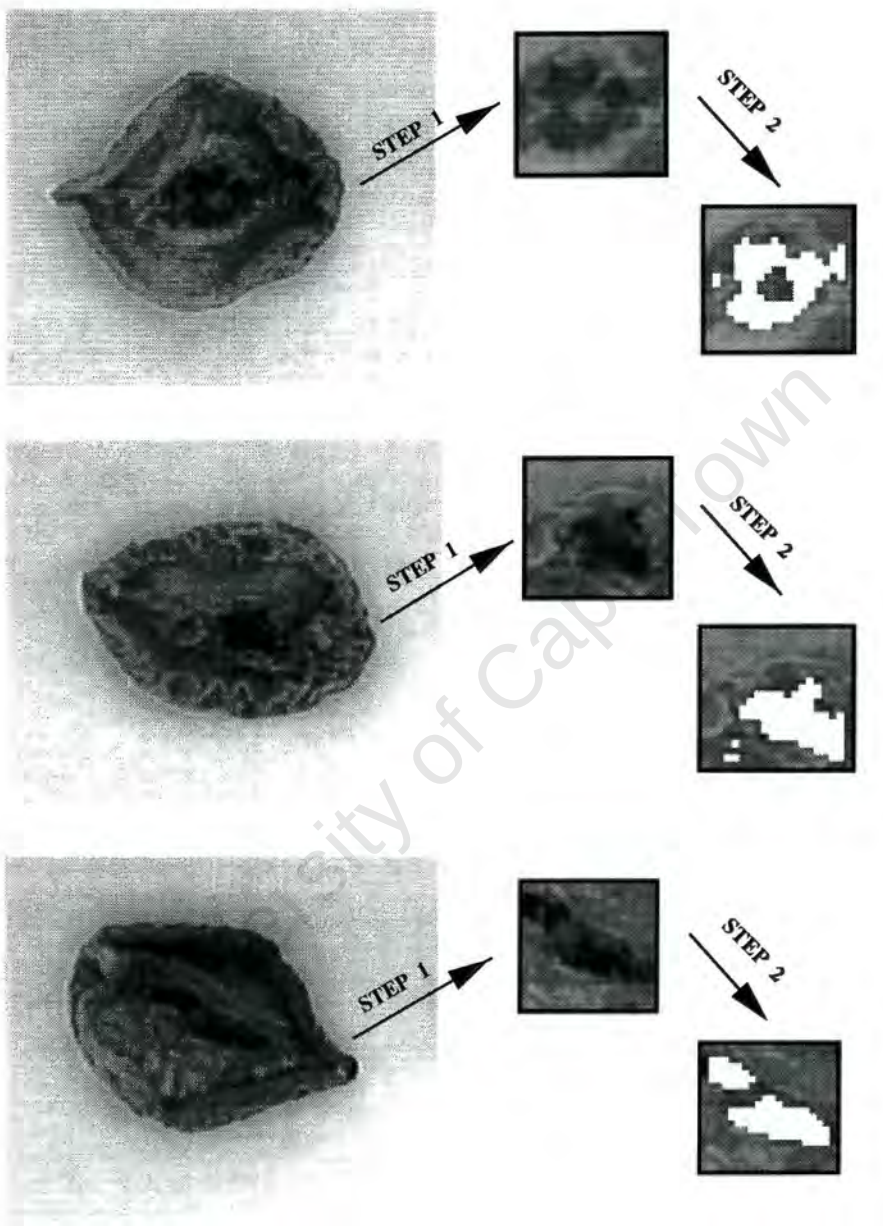

STEP 1 : A region about the centre-of-mass is isolated

STEP 2 : The Top- Hat Transform removes the dark core area

Figure 5.10: The results of the Top-Hat Transform on image Red band values

dried peach. Simple statistical tests (see Appendix) of the fruit image indicate whether or not the image is a representation of the 'front' or 'back' of the peach. This ensures that no unnecessary processing is performed.

The above Top-Hat transform segmentation procedure can be used to isolate any known dark area on a fruit: tests on dried bananas showed that the dark centre line of a dried banana was isolated very effectively from the remainder of the banana flesh. In this case the window was reset to encompass the length of the fruit, rather than a square area about the centre-of-mass of the fruit.

In this Chapter basic theories and concepts of binary mathematical morphology were introduced. The core region of a dried peach was successfully segmented using the Top-Hat transform operation; both the operation and the segmentation process were described.

The last two chapters have dealt with the three segmentation tasks that were necessary to ensure a successful colour analysis of selected dried fruit types : the dried fruit was isolated within the image, dark coloured peach core areas were removed from the dried peach flesh, and black blemishes on dried pears were isolated and *labeled*. In Chapter 3, the RGB colour space was selected to best describe the colour grades of dried fruit. The remaining problem is that of finding a classification system that will automatically grade each fruit according to its $[R, G]$ feature vector. A number of relevant classification methods are discussed in Chapter 6.

Section 5.5: Application of the Top-Hat Transform to remove the core area from a dried peach

Chapter 6

Discrimination and Classification of Multivariate Data

6.1 Introduction

In Chapter 2 the three basic steps in a general pattern recognition problem were described as : (1) the 'sensing problem' where a suitable *measurement vector* was set up from the input source, (2) the extraction and possible reduction in dimensionality of a suitable *feature vector* which accurately characterised the patterns, and finally (3) the determination of a decision procedure or classification scheme which would automatically determine to which category or class a given unit belonged, where a classification (or making a decision on the class assignment of the input pattern) is based on the measurements taken from the selected features [27].

The image source in this dried fruit system, as described in Chapter 2, is the RGB CCD camera, which produces measurements of the red, green and blue light reflected from the image scene; these measurements constitute the $[R, G, B]$ *measurement vector.* In Chapter 3 various colour models were examined in order to determine the most appropriate feature vector to represent the colour in the dried fruit. The RGB colour model was selected; the optimal dimensionality of the $[R, G, B]$ feature vector was then reduced to $[R, G]$. The $[R, G]$ feature vectors are those that will define the feature space of the classification system, and are the vectors which, when submitted to a classifier, must be allocated by the classifier to the correct class. *Multivariate data* arise whenever an investigation of a physical phenomenon is described using a number $\nu \geq 1$ of variables or characters; the values of these variables are all recorded for each distinct item or experimental trial in a set of n items or trials [37]. In the dried fruit grading system, the R and G variables constitute the multivariate data of

the system; each new dried fruit has its own average R, G values which describe its individual colour information.

In order to examine data sets, it is often easier to compare or describe them using corresponding summary statistics. Three of the most commonly used concepts in summary statistics of multivariate data sets, are the data set **average**, **variance** and **covariance**; these are detailed in Section 6.2.

There are essentially two types of classification systems: the *statistical* or *probabilistic* pattern recognition system, and the *deterministic* or *non-statistical* pattern recognition system [22]. Haralick [38] describes the *deterministic* system as one where, for some class a, if the probability of an element being in class a is 1, then for all other classes the probability of that element is 0; decision rules that are not *deterministic* are called *probabilistic*. Both classification types aim to find the optimum decision function or rule, which will correctly classify or match any feature vector to its correct category or class. The difference between the two approaches lies in the availability of the *a priori* probabilities, or *training sets*, for the respective data or data classes.

If *a priori* belief or knowledge about the patterns to be recognised, is available in a probabilistic form, then the selection of the optimal decision function can be generated based on *Bayesian theory*. Then the *Bayesian probability of error* is an effective measure of a decision rule's usefulness [23]. If *a priori* knowledge is not available in that form, but training sets with known item classification are available, as is the case in this system [22], then the decision rule options include Linear Discriminant Functions or Heuristic Non-linear boundary fitting to the data classes. If no *a priori* knowledge, or suitable training sets are available, clustering is used. Classifiers that are developed according to *Bayes decision rules*, are however, optimal [38] in the sense that they incorporate prior experience or belief as additional pseudo-data.

Multivariate data statistics and distance measurement are important elements in pattern recognition classifiers, and are therefore introduced in Section 6.2. Probability density functions, and in particular the normal density function which is often used to model and describe class data distributions in pattern recognition applications [37] are described in Section 6.3.

The Bayesian conditional loss classification approach is described in Section 6.4. In Section 6.5 Linear Discriminant functions are introduced; more expansive discussions on the detail and applications of Linear discriminant functions can be found in [54], [27], [41], [22]. An overview of Clustering is given in Section 6.6.

In a dried fruit grading system, *training sets* will be available which will train the classifier; these independent training sets will be used to set up the parameters and therefore the decision rules for the colour grading classification scheme. These parameters will then be used by the classifier to allocate any sample unit of dried fruit to the correct class or grade. The details

of this *learning* procedure are in Section 6.7.

In order to gauge the success of the results of a classifier method, the *error* must be calculated. Error analysis of the chosen classifier is discussed in Section 6.8. The **Jeffries-Matusita distance measure**, a test on the separability and therefore the success, of the classes set up by the classifier, was proposed by Niblack [51]. This test is also detailed in Section 6.8. A final summary of the procedures adopted to classify dried fruit are given in Section 6.9.

6.2 Multivariate data and Descriptive Statistics

Much of the information contained in the multivariate data can be assessed by calculating and using certain summary numbers from the multivariate data sets; these numbers are known as *descriptive statistics* and involve the concepts of **sample mean**, the **sample variance** and the **sample covariance**.

6.2.1 Descriptive Statistics

The Sample Mean

The *arithmetic average* or *mean* of n measurements $(x_{11}, x_{12}, ..., x_{1n})$ of a variable X_1, denoted by $\overline{x_1}$ is given by :

$$\overline{x_1} = \frac{1}{n}\sum_{j=1}^{n} x_{1j}$$

If the n measurements represent a subset of the full set of measurements that might have been observed, then $\overline{x_1}$ is also called the *sample mean* for the first variable x_1. In the case of multivariate data where the number of variables $\nu > 1$, the *sample mean* can be computed from n measurements on each of the ν variables, so that in general there will be ν *sample means* in a sample mean vector with elements :

$$\overline{x_i} = \frac{1}{n}\sum_{j=1}^{n} x_{ij} \qquad\qquad i = 1, 2, ..., \nu$$

The Sample Variance

A measure of spread or variation in the data values of a variable X_1 is provided by the *sample variance* ${s_1}^2$. This statistic is defined for n measurements on the first variable x_1 as :

$${s_1}^2 = \frac{1}{n}\sum_{j=1}^{n} (x_{1j} - \overline{x_1})^2$$

where $\overline{x_1}$ is the sample mean. In some contexts divisor $(n-1)$ is used in place of n.

The Sample Covariance

Consider n pairs of measurements on each of variables x_1 and x_2 :

$$\begin{bmatrix} x_{11} \\ x_{21} \end{bmatrix}, \begin{bmatrix} x_{12} \\ x_{22} \end{bmatrix}, \cdots, \begin{bmatrix} x_{1n} \\ x_{2n} \end{bmatrix}$$

x_{1j} and x_{2j} are observed on the j^{th} experimental item where $(j = 1, 2, ..., n)$. A measure of **linear association** between the measurements of variables X_1 and X_2 is provided by the *sample covariance* s_{12} :

$$s_{12} = \frac{1}{n}\sum_{j=1}^{n}(x_{1j} - \overline{x_1})(x_{2j} - \overline{x_2})$$

or the average product of deviations from their respective means. Again in some contexts divisor $(n-1)$ is used.

Note: *Sample* mean and covariance values must not be confused with *expected* mean and covariance values: *expected* values include consideration of a probability $p_i(x_i)$ for x_i, where X_i is a discrete random variable, or probability element $f(x_i)dx_i$ where X_i is a continuous random variable.

6.2.2 Multivariate Distance measurement

Many multivariate classification techniques are based upon the simple concept of distance measurement. The straight-line or *Euclidean distance* is the most familiar [37], but often a more flexible *generalised distance* measure is appropriate. The structure of the data sample measurements justify the choice of distance measurement, and dictate the geometry for the low dimensional representation of the ν-dimensional data [37].

If we consider point $P = (x_1, x_2)$ in a 2-dimensional plane, the straight-line or *Euclidean distance* $d(O, P)$ from P to the origin $O = (0, 0)$ is, according to Pythagorean theorem :

$$d(O, P) = \sqrt{{x_1}^2 + {x_2}^2}$$

In general the *Euclidean distance* from a point $P = (x_1, x_2, ..., x_p)$ to a point $Q = (y_1, y_2, ..., y_p)$ is :

$$\begin{aligned} d(P, Q) &= \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_p - y_p)^2} \\ &= \sqrt{(\mathbf{x} - \mathbf{y})^T(\mathbf{x} - \mathbf{y})} \end{aligned} \tag{6.1}$$

where T implies the transpose operator on the preceding vector or matrix.

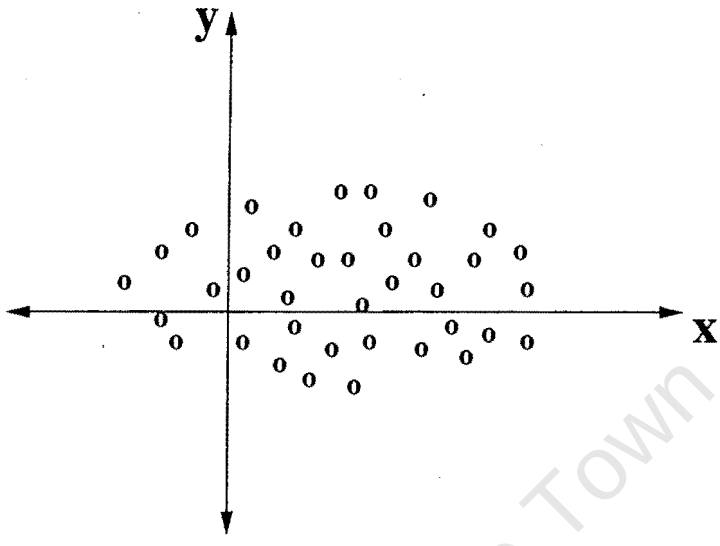

Figure 6.1: A scatterplot showing x variations larger than y variations

In *Euclidean distance* measurement, each coordinate contributes equally to the calculation of the distance; this balance is not always satisfactory for statistical purposes. In some cases the variable measurements are subject to different rather than similar extents of fluctuations. What this means is that two variables X and Y may not undergo similar variations, as is assumed in the *Euclidean distance* case; changes in the X variable measurements, for example, may tend to be obviously greater than changes in the Y variable. The *scatterplot* of Figure 6.1 shows such an example. In these cases it would be preferable to weight those coordinate variables such as X, that fluctuate greatly, less heavily than variables like Y. These distance measures that account for differences in variable variation are referred to as *statistical distances*. For variable measurements which **do not vary independently**, the *statistical distance* measure between points P and Q is :

$$d_s(P,Q) = \sqrt{(\mathbf{x}-\mathbf{y})^T \mathbf{A} (\mathbf{x}-\mathbf{y})} \tag{6.2}$$

Ordinarily $\mathbf{A} = \mathbf{S}^{-1}$ where the sample covariance matrix $\mathbf{S}$ contains the sample variances and covariances [37] of Section 6.2.1.

Without prior knowledge of the data groups or classes, the sample mean, variance and covariance can only be computed for the aggregated data set; in these cases the *statistical distance* measure will not be useful for classification.

6.2.3 Relevance of statistical distance measure to R,G values

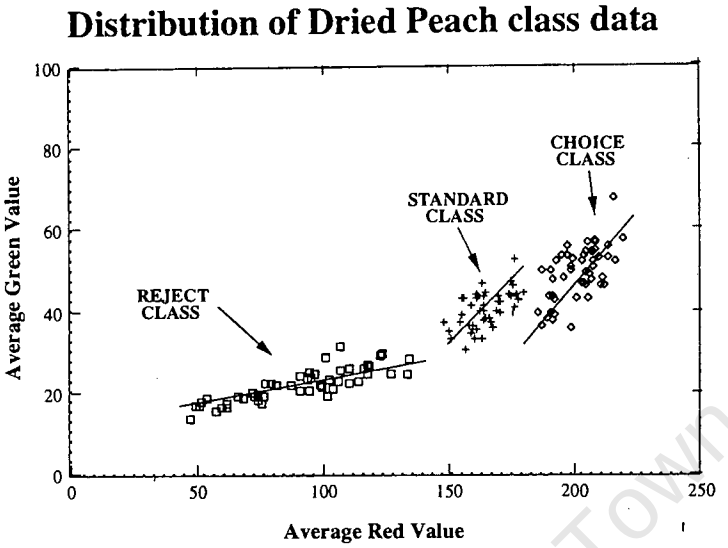

Figure 6.2: A scatterplot showing class differences in position and in relative variations within X and Y

Figure 6.2 is a scatterplot which describes three classes of dried peach: Choice, Standard and Reject, in a 2-dimensional R, G feature space. Each point in the scatterplot describes the average R, G values for one individual dried peach (with the core area removed). In each class, the data varies more in distance along one axis than the other; the data are elliptical in nature. For this reason the *statistical distance* $d_s(P, Q)$ of Equation 6.2 would serve as a more suitable distance measure technique than the Euclidean measure.

This distance representation of the data distribution facilitates the use of a probability density function to describe each dried fruit class distribution. Gaussian or normal probability density functions are based on a statistical square distance measure of the samples.

If each dried fruit class could be modelled by a normal distribution, then the dried fruit *training set* of a class would be used to approximate the parameters of the corresponding normal probability distribution.

6.3 The Multivariate Normal Distribution

The general form of a family of important probability density functions is given by the relation[1]

$$p(\mathbf{x}) = K_\nu \, |\mathbf{\Sigma}^{-1}|^{\frac{1}{2}} \, f[(\mathbf{x}-\mu)^T \;\; \mathbf{\Sigma}^{-1}(\mathbf{x}-\mu)] \tag{6.3}$$

where

- K_ν is a normalising constant
- $\mathbf{\Sigma}$ is the true covariance matrix which is real, symmetric and positive definite
- μ is the true mean vector
- ν is the dimensionality of $\mathbf{x}$.

The sampling distributions of many multivariate statistics are approximately normal, regardless of the form of the parent population, because of the **Central-Limit Theorem** [2]. The multivariate normal density is a member of this family of density functions, and is easily understood if it is generalised from the univariate normal density function to ($\nu \geq 2$) dimensions.

6.3.1 Generalisation of the univariate normal density function to the multivariate normal density function

The univariate normal distribution with mean μ, and variance σ^2 has the probability density function :

$$\begin{aligned} f(x) &= \frac{1}{\sqrt{2\pi\sigma^2}} e^a \\ \text{where } a &= -\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2 \qquad\qquad -\infty < x < \infty \end{aligned} \tag{6.4}$$

The term $\left(\frac{x-\mu}{\sigma}\right)^2$ in the exponent of the multivariate normal density function can be written :

$$\left(\frac{x-\mu}{\sigma}\right)^2 = (x-\mu)(\sigma^2)^{-1}(x-\mu) \tag{6.5}$$

[1] To take class membership into account the notation for $p(\mathbf{x})$ would be :

$$q(\mathbf{x}/w_i) = K_\nu \, |\mathbf{\Sigma}_i^{-1}|^{\frac{1}{2}} \, f[\, (\mathbf{x}-\mu_i)^T \;\; \mathbf{\Sigma}_i^{-1}(\mathbf{x}-\mu_i) \,]$$

[2] Let $X_1, X_2, ..., X_n$ be independent ν-dimensional observations from any population with mean μ and finite covariance matrix σ. Then $\sqrt{n}(\overline{X}-\mu)$ has an approximate normal distribution $N_\nu(0, \Sigma)$ for large sample sizes. Here n should be large relative to ν [37].

This term measures the squared distance from x to μ in standard deviation units. For a vector $\mathbf{X}$ of observations on several variables, the generalised squared distance becomes :

$$(\mathbf{x}-\mu)^T\mathbf{\Sigma}^{-1}(\mathbf{x}-\mu) \tag{6.6}$$

where the $[\nu \times \nu]$ matrix σ represents the covariance matrix. For the multivariate form, the univariate normalising constant $(2\pi)^{\frac{-1}{2}}(\sigma^2)^{\frac{-1}{2}}$ is changed to $(2\pi)^{\frac{-\nu}{2}}\mid \sigma \mid^{\frac{-1}{2}}$ to ensure that the volume under the surface of the multivariate density function is unity for any ν. The multivariate normal probability distribution $p(\mathbf{x})$ is then :

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{\nu}{2}}\ |\mathbf{C}|^{\frac{1}{2}}}\ e^a \tag{6.7}$$

$$a = [-\frac{1}{2}(\mathbf{x}-\mu)^T\ \mathbf{\Sigma}^{-1}(\mathbf{x}-\mu)]$$

which belongs to the class of density functions defined by Equation 6.3 where

$$K_\nu = (2\ \pi)^{\frac{-\nu}{2}}$$

$$f[\cdot] = exp\ [\ -\tfrac{1}{2}\ (\mathbf{x}-\mathbf{m})^T\ \ \mathbf{C}^{-1}(\mathbf{x}-\mathbf{m})\ ]$$

$$\hat{\mu} = \mathbf{m} = \simeq \tfrac{1}{n}\ \textstyle\sum_{j=1}^{n}\ \mathbf{x}_j, \quad \text{where n is the number of samples}$$

$$\hat{\mathbf{\Sigma}} = \mathbf{C} = \simeq \tfrac{1}{n}\ \textstyle\sum_{j=1}^{n}\ (\mathbf{x}_j\mathbf{x}_j^T - \mathbf{m}\mathbf{m}^T)$$

The above estimate of these two parameters are derived by the Maximum Likelihood (ML) criterion for Parameter Estimation. Under the ML criterion the sample mean and sample covariance matrix are derived from the data *training set.* The parameter estimates for each class, the mean $\mathbf{m}$ and covariance matrix $\mathbf{C}$ must be obtained. A distinct μ and $\mathbf{\Sigma}$ applies within each class.

6.3.2 Dried fruit class distribution described by normal probability densities

If a normal probability distribution is used to model each of the dried fruit data sets, then for each class of fruit w_i, a data-based normal distribution $p_i(\mathbf{x})$ must be generated in the R, G feature space :

$$p_i(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{2}{2}}\ |\mathbf{C_i}|^{\frac{1}{2}}}\ e^a$$

$$\text{where} \quad a = [-\frac{1}{2}\ (\mathbf{x} - \mathbf{m_i})^T\ \mathbf{C_i}^{-1}(\mathbf{x} - \mathbf{m_i})]$$

For this 2 band problem, $\mathbf{m_i}$ is a 2-dimensional vector containing both the **R sample mean** and the **G sample mean** and $\mathbf{C_i}$ is a (2 x 2) **sample covariance matrix**.

The above discussion on probability distributions is an informal introduction to the topic, the aim being to highlight the areas relevant to understanding the design of a Classifier based on normal probability distributions. More formal and thorough explanations can be found in [22], [41], [30] and most other text covering statistics in Pattern Recognition.

6.4 Bayes Classification

Bayes decision theory is a statistical approach to pattern classification. This approach is based on the assumption that the decision problem is posed in probabilistic terms, and that all of the relevant *a priori* probabilities are known [39]. Only the basic Bayesian formula is described in this section; a clear and simple explanation of the basis and derivation of the Bayes decision theory is described in Haralick [38], and a more thorough and detailed discussion of Bayes classification is found in Fukunaga [23] and Young [49].

The Bayes formula for upgrading the probabilities of membership of classes i, taking the items data $\mathbf{x}$ into account is :

$$p(w_i|\mathbf{x}) \quad = \quad \frac{p(w_i)\ p(\mathbf{x}|w_i)}{p(\mathbf{x})}$$

where

$$p(\mathbf{x}) \quad = \quad \sum_{i=1}^{M} p(\mathbf{x}|w_i)\ p(w_i)$$

operates as a normalisation factor and

- the possible pattern classes are w_i for $i = 1, .., M$.
- $p(w_i)$ is the known or assumed *a priori* probability of any element $\mathbf{x}$ belonging to class w_i.
- $p(\mathbf{x}|w_i)$ is the likelihood function of class w_i, or the conditional probability of observing measurement $\mathbf{x}$ given that the class is w_i. A common choice of functional representation of $p(\mathbf{x}|w_i)$ is as a Gaussian or Normal probability distribution in which the function

parameters are estimated as the observed class i **mean** and the **covariance matrix** (ie the sample estimates for class i) [51].

- $p(w_i|\mathbf{x})$ is the *a posteriori* probability of $\mathbf{x}$ being in class w_i, that is, observing the value of $\mathbf{x}$ changes the *a priori* probability $p(w_i)$ to the *a posteriori* probability $p(w_i|\mathbf{x})$.
- $p(\mathbf{x})$ is the likelihood function of the mixture of all classes.

The function of a pattern classifier is to perform the classification task minimising the penalty associated with the misrecognition of items. What this means is that if the classifier decides that $\mathbf{x}$ belongs in class w_j, when it actually belongs in w_i, it incurs a positive loss $\mathbf{L_{ij}}$ where $\mathbf{L_{ii}} = 0$. Since pattern $\mathbf{x}$ may belong to any of the M classes under consideration, the expected loss in assigning $\mathbf{x}$ to class w_j, referred to as the *conditional expected loss* for misclassifying $\mathbf{x}$ into w_j, is

$$r_j(\mathbf{x}) = \sum_{i=1}^{M} L_{ij}\ p(w_i|\mathbf{x})$$

Suppose we define a region R_j so that if $\mathbf{x}$ is in R_j we assign $\mathbf{x}$ to w_j $(j = 1, 2, .., M)$. Then the average unconditional loss for misclassifying $\mathbf{x}$ into w_j is

$$\mathbf{L}_j = \int_{R_j} r_j(\mathbf{x})\ p(\mathbf{x})dx$$

and the total loss is

$$\mathbf{L} = \sum_j L_j\ p(w_j)$$

To minimise $\mathbf{L}$ we select the appropriate regions R_j so that each $\mathbf{L}_j$ is minimal :

$$\begin{aligned} \min_j \mathbf{L}_j &= \min_j \int_{R_j} \sum_i \mathbf{L}_{ij}\ p(w_i|\mathbf{x})p(\mathbf{x})dx \\ &= \min_j \int_{R_j} \sum_i \mathbf{L}_{ij}\ p(w_i|\mathbf{x})p(w_i)dx \\ &= \sum_i \min_j \int_{R_j} \mathbf{L}_{ij}\ p(w_i|\mathbf{x})p(w_i)dx \end{aligned} \tag{6.8}$$

Thus the *Bayes rule* is to assign $\mathbf{x}$ to that w_j for which $\sum_i \mathbf{L}_{ij}\ p(w_i|\mathbf{x})p(w_i)$ is minimal, and R_j is the region over which the minimum is achieved by class w_j.

In order to determine the Bayes decision rule of the Bayes classifier, the *a priori* probability functions $p(w_i)$ must be known. When a situation exists in which the prior probabilities cannot be known, and it is unreasonable to assume that they are equal, the Bayes decision rule or classifier cannot be applied [38].

6.5 Discriminant Functions

6.5.1 Linear Discriminant Functions

Points that occur within a small Euclidean distance of one another in the pattern space, and cluster into a distinct group, normally indicate a specific pattern class. If boundaries can be set up between such groups of points, that divide the clusters from one another, then the classes can be uniquely specified. If it is assumed that the machine will be designed to recognise M different pattern classes, denoted by $(w_1, w_2, ...w_M)$, the pattern or feature space can then be considered as consisting of M regions, each of which encloses the pattern points of a class. The recognition problem can then be viewed as that of generating the decision boundaries which separate the M classes. The decision boundaries can be defined by *decision* or *discriminant* functions $d_1(\mathbf{x}), ...d_M(\mathbf{x})$. The discriminant functions can be generated in a variety of ways. The functions can be linear functions, piecewise linear functions, or non-linear functions.

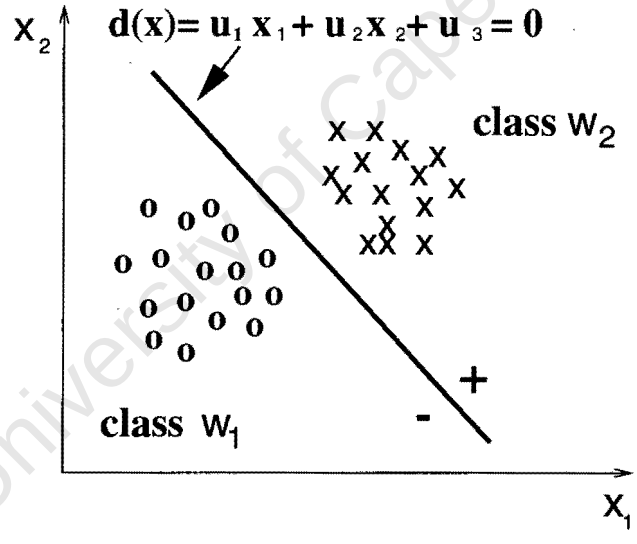

Figure 6.3: A simple Discriminant function for two pattern classes

Tou [22] introduces the concept of linear discriminant functions by way of a simple example : A feature vector $\mathbf{x} = (x_1, x_2)$ consists of two measurements, x_1 and x_2. The two pattern classes into which the pattern or feature space must be divided are w_1 and w_2. As can be seen in figure 6.3 this particular pattern population can be separated into two distinct groups by a boundary line or rather by the values of a *decision function* $d(\mathbf{x})$. The equation for the surface separating the 2 pattern classes is obtained by letting the decision function equal

zero :

$$d(\mathbf{x}) = u_1 x_1 + u_2 x_2 + u_3 = 0$$

where u_1, u_2, u_3 are parameters
x_1, x_2 are the co-ordinate variables.

In Figure 6.3 any pattern $\mathbf{x}$ belonging to class w_2 will yield a positive scalar value when substituted into $d(\mathbf{x})$, hence $d(\mathbf{x}) > 0$. Similarly, any pattern belonging to class w_1 would yield a negative value upon substitution into $d(\mathbf{x})$, hence $d(\mathbf{x}) < 0$. The discriminant function $d(\mathbf{x})$ could be used within a class classifier, since any pattern $\mathbf{x}$ of unknown classification would belong to :

1. class w_1 if $d(\mathbf{x}) < 0$ or
2. class w_2 if $d(\mathbf{x}) > 0$

If the pattern $\mathbf{x}$, however, fell upon the decision boundary $d(\mathbf{x}) = 0$, then the condition is indeterminate, and the pattern would not be uniquely allocated to either class. Such patterns might be randomly assigned to one of the classes.

General Form of the n-dimensional Linear Decision Function

The above 2-dimensional linear discriminant function when generalised to the n-dimensional case is of the form [22] :

$$\begin{aligned} d(\mathbf{x}) &= w_1 x_1 + w_2 x_2 + ...w_n x_n + w_{n+1} \\ &= \mathbf{w}_0^T \mathbf{x} + w_{n+1} \end{aligned}$$

$$\begin{aligned} \text{where } \mathbf{w}_0 &= (w_1, w_2, ...w_n)^T \\ \mathbf{x} &= (x_1, x_2, ...x_n) \end{aligned}$$

$\mathbf{w}_0$ is referred to as the *weight* or *parameter* vector, $\mathbf{x}$ is the *pattern vector*, and T implies the transpose vector. It is a widely accepted convention to express $d(\mathbf{x})$ in the form [22] :

$$d(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$$

$$\begin{aligned} \text{where } \mathbf{w} &= (w_1, w_2, ...w_n, w_{n+1})^T \\ \mathbf{x} &= (x_1, x_2, ...x_n, 1)^T \end{aligned}$$

In this form $\mathbf{w}$ is the *augmented weight vector* which now includes the constant, w_{n+1}, and $\mathbf{x}$

is the *augmented pattern vector* which now includes a constant, 1.

The n-dimensional case of a linear discriminant classifier is then :

$$d_i(\mathbf{x}) = \mathbf{w}_i{}^T\mathbf{x} = \begin{cases} > 0 & \text{if } \mathbf{x} \in w_i \\ < 0 & \text{otherwise} \end{cases}$$

where $\mathbf{w}_i = (w_1, w_2, ...w_n, w_{n+1})^T$ is the *augmented weight vector* associated with the i^{th} discriminant function and $i = 1, 2, ..., M$.

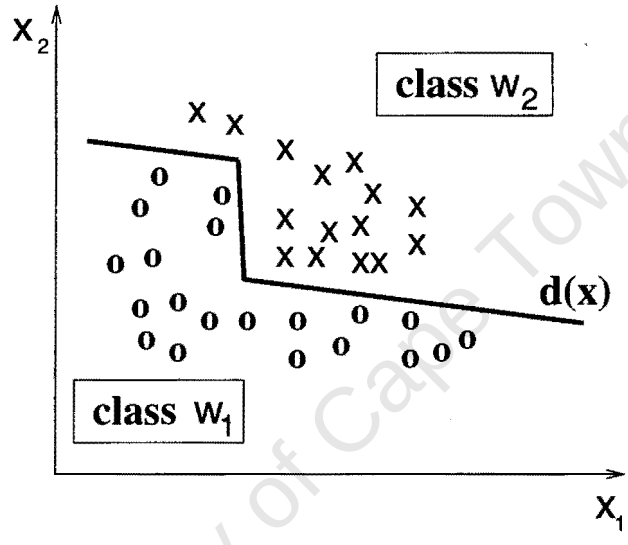

Figure 6.4: A Piecewise Linear Discriminant function

A Piecewise Linear function is a function that is linear over subregions of the feature space. These functions give piecewise linear boundaries between classes. Figure 6.4 is an example of a Piecewise Linear function. As with simple linear discriminant functions, the decision rules should correctly allocate a pattern vector to its class. Further detail on the application of piecewise linear decision functions, as well as various examples are described in both [41] and [22].

The success of a pattern classification scheme which uses Discriminant functions depends upon two factors [22] :

1. the form of $d(\mathbf{x})$
2. ability to determine the discriminant coefficients

The first problem is related to the geometrical properties of the pattern classes under consideration [22]. In the second problem certain techniques can be used to locate an optimum

solution. If there exists little, if any knowledge about the patterns to be recognised, pattern recognition machines are best designed using a *training* or *learning* procedure with the aid of *training sets* [22]. In these cases initial decision functions can be assumed, and through a sequence of training steps, these decision functions can be made to approach satisfactory forms. Several adaptive and training techniques that will train the system to achieve the optimum weight vector, **w** are found in [22]. Another important approach, is the determination of a discriminant function using the assumption that the data in the classes follow normal distributions. In this case the discriminant function parameters are based entirely on the parameters of the normal distributions (the **mean** and **covariance matrix**), which in turn are estimated from the *training sets.* Discriminant functions that separate data sets that have unequal variances or covariances are called *Quadratic Discriminant functions.* In Section 6.5.2, 6.5.3 and 6.5.4, three types of discriminant function, which segregate dried fruit classes, are detailed: the first is an *elliptical discriminant function* which bounds each class within an ellipse, the second is a purely *linear discriminant function,* and the third is a *quadratic discriminant function* which will segregate the normal density distributions that describe each dried fruit class. Each method will be applied to the full data sets for three classes of dried peaches; each of these classes consist of 51 peaches, all of which were hand-picked by trained staff of South African Dried Fruit [34].

6.5.2 Non-linear Discriminant Functions

In certain circumstances pattern groups which cluster into distinctly shaped groups can be bounded by non-linear decision functions: simple examples would include circular or elliptical boundaries.

Fitting an ellipse to each training set or class of dried fruit

In the (R, G) feature space selected, most observed class data distributions tended to be elliptical in shape. By applying a Chi-square linear fit to the data in the training set of each class, the major axis of a bounding ellipse was found. Each data point was then translated onto a 0^o axis; the outlying points on this 0^o axis were then labelled the two foci of the ellipse. An ellipse that included every data point of the class training set was then generated. In Figure 6.5 three of the four classes of dried peaches are plotted in (R, G) space. Major axes and bounding ellipses have been added by hand in order to show the elliptical nature of the data. Classes of data points that displayed more of a circular or non-elliptical geometry would however, not be suitably defined using this discrimination method because the Chi-square linear fit would not be successful in fitting a major axis to such data distributions. If this discrimination method was adopted, further consideration of those observations which fell in the feature space where two ellipses overlapped, would be necessary. These particular

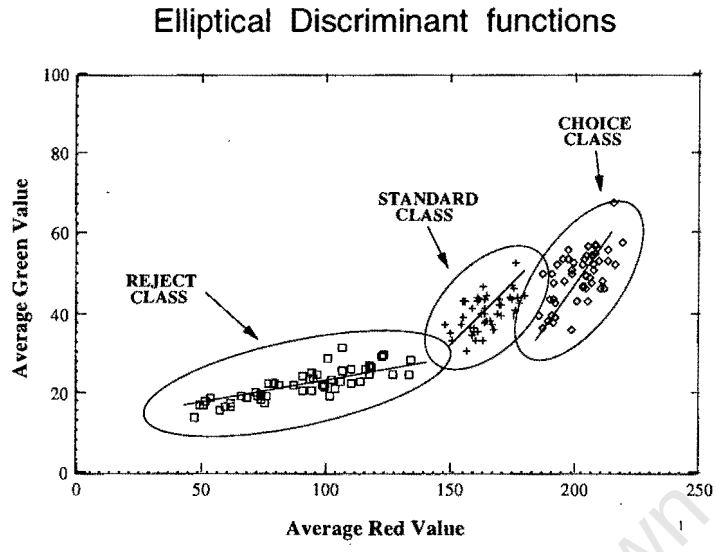

Figure 6.5: Elliptical decision boundaries to separate classes of dried peaches

observations would have to be classified into only one of the two classes, using other factors or weighting techniques.

6.5.3 Linear Discriminant functions to classify Dried Fruit

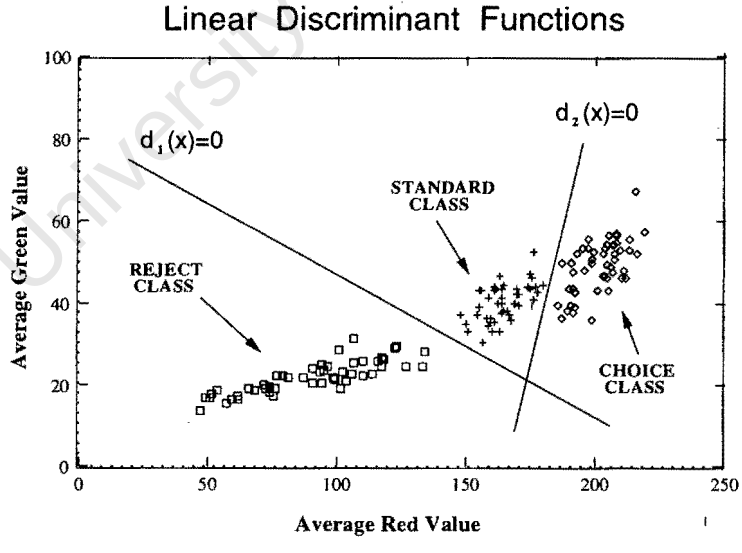

Figure 6.6: Linear decision boundaries to separate classes of dried peaches

Although certain dried fruit colour classes were expected to be very close to one another, none should overlap - i.e. a peach cannot be both a choice grade as well as a standard grade peach. The groups of $[R, G]$ feature vectors which defined each class were expected to cluster into

distinct classes. A minimum of two linear decision functions would determine the boundaries between three classes of dried peach. A 3-class linear discriminant system for the classes of Choice, Standard, and Reject dried peaches is illustrated graphically in Figure 6.6. The full *labeled sets* of 50 dried peaches are described in Figure 6.6. If linear decision functions were used to classify the dried fruit, the best weight vector $\mathbf{w_o}$, would have to be found.

Linear decision functions, however, are suitable for data sets where the observed covariance matrices are equal across classes. As can be seen in Figure 6.6, which describes the full data set for three dried peach classes, it is unlikely that the dried fruit classes have equal variance structures. Discriminant functions that separate data sets that have unequal covariance matrices are called *Quadratic Discriminant functions.* It is clear therefore, that classification of dried fruit classes, where each class is described by a normal density distribution, would require *quadratic discrimination.*

6.5.4 Quadratic Discriminant functions to classify Dried Fruit

Pattern recognition algorithms are often classed as *parametric* or *non-parametric.* A parametric approach to pattern classification derives the decision function in part from a class of probability densities that are defined by a relatively small number of parameters. In *parametric methods* it is common to assume that each pattern class arises from a multivariate normal distribution, where the parameters are the data set *mean vector* and the *covariance matrix.* On the other hand, although non-parametric methods often use parameterised decision functions, no conventional form of probability distribution is assumed [54].

If each dried fruit class was defined by a normal distribution, which by definition has no finite borders, any dried fruit observation would have a probability of being in all of the four classes - albeit an extremely small probability of occurrence in some classes. The sum of the probability densities multiplied by the corresponding class probabilities defines the combined probability surface

$$f(\mathbf{x}) = \sum_{i=1}^{M} p(\mathbf{x}|w_i).p(w_i)$$

as in Figure 6.7.

In the training run the parameters of class mean vector and covariance matrix of the normal probability distribution for each class would have to be defined. In this approach the linear decision boundary between two classes would be the 'shadow line' or projection onto the X-Y surface, from the interconnecting surface of two normal distributions weighted by the corresponding underlying class probabilities; at this interconnecting surface the probability that an observation is from either class is $\frac{1}{2}$.

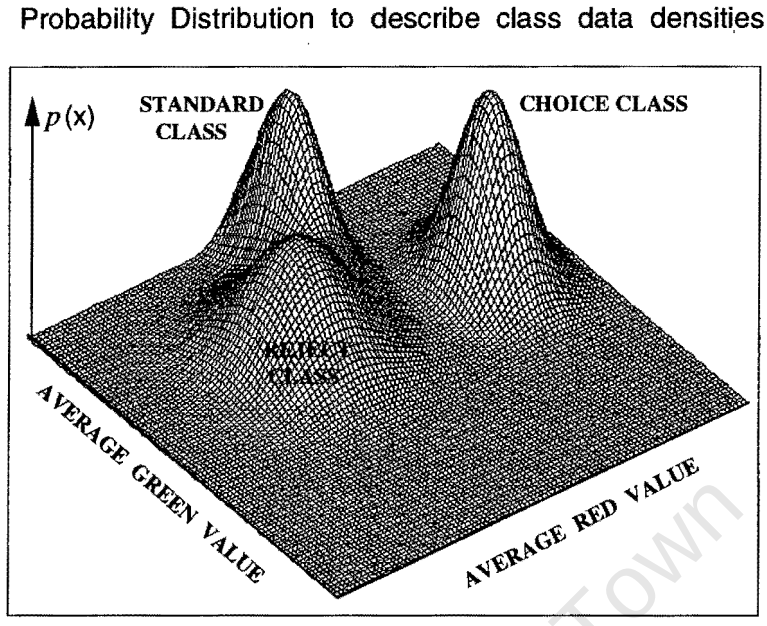

Figure 6.7: Mixture of Distinct normal Probability distributions representing each class in a combined probability distribution for **x** of an unknown class

6.5.5 Optimal Classification Rules

A good classification procedure should result in few misclassifications. Over and above the basic discriminant rule, there are two additional features that an 'optimal' classification rule should possess [37]: one feature takes the *a priori* class probabilities $p(w_i)$ and the probability of occurrence into account; the other takes the *cost of misclassification* into account. The *costs of misclassification* can be defined by a cost matrix such as the following example :

True Population	Classify as: Choice	Reject
Choice	0	c(R | C)
Reject	c(C | R)	0

The **costs** are :

1. Zero for correct classification
2. c(C | R) when an observation from Reject is incorrectly classified as Choice
3. c(R | C) when a Choice observation is incorrectly classified as Reject.

Johnson [37] suggests the following classification rule for assigning $\mathbf{x}$ to one of classes R1 and R2 :

$$R1 : \frac{f_1(\mathbf{x})}{f_2(\mathbf{x})} \geq [\frac{c(1|2)}{c(2|1)}][\frac{p_2}{p_1}] \tag{6.9}$$

$$[\text{Density ratio}] \geq [\text{cost ratio}]\ [\text{prior probability ratio}]$$

$$R2 : \frac{f_1(\mathbf{x})}{f_2(\mathbf{x})} < [\frac{c(1|2)}{c(2|1)}][\frac{p_2}{p_1}] \tag{6.10}$$

$$[\text{Density ratio}] < [\text{cost ratio}]\ [\text{prior probability ratio}]$$

where $f_1(\mathbf{x})$, $f_2(\mathbf{x})$ are the class density distributions, and p_1, p_2 are the probability of occurrences.

No *a priori* class probabilities are available for dried fruit. The *a priori* probabilities for each group are taken to be equal; what this means is that a dried fruit $\mathbf{x}$ has the same probability of being in Grade 1, as it has for any of the other classes, which is an acceptable assumption under specified test set conditions. The costs of misclassification would need to be determined by the respective company; for test purposes of this dried fruit classifier, costs of misclassification have been assumed equal. Johnsons [37] classification rule of Equation 6.10, under the above two assumptions becomes :

$$R1 : \frac{f_1(\mathbf{x})}{f_2(\mathbf{x})} \geq 1$$

$$R2 : \frac{f_1(\mathbf{x})}{f_2(\mathbf{x})} < 1 \tag{6.11}$$

for a 2-class classifier. This means that a pattern $\mathbf{x}$ has a degree of membership in every class, but it is a member of the class to which it has the highest degree of membership, that is

$$\mathbf{x} \in \text{class } i \qquad \text{if } p_i(\mathbf{x}) > p_j(\mathbf{x}) \qquad \text{for } j \neq i$$

6.6 Clustering

In Supervised learning the system is trained to classify a known set of patterns using *a priori* information or training sets, thereafter the system can freely classify other patterns. In Non-supervised learning, little, if any *a priori* knowledge about the data structure, or the number of groups is known. In these situations one would need to study the inter-relationships among the data points in order to make a preliminary assessment of their structure. This strategy would fall under the domain of exploratory data analysis; **cluster analysis** is one tool of exploratory data analysis that attempts to assess the interaction among patterns by organising the patterns into groups or clusters, such that patterns within a cluster are

more similar to each other than are patterns belonging to different clusters [49]. Grouping is done on the basis of similarity or distances (dissimilarities) [37]. The similarity measure (or dissimilarity measure) is usually given in numerical form to indicate the degree of natural association or resemblance between patterns in a group or between pattern groups. These measures would include the statistical distance, or more commonly the Euclidean distance of Section 6.2. A thorough mathematical description of distance and similarity functions can be found in [12].

Given N feature or pattern vectors $\mathbf{x}_1, \mathbf{x}_2, ..\mathbf{x}_N$ contained in the feature space $\mathbf{S}$, the process of clustering can be described as seeking the regions $\mathbf{S}_1, \mathbf{S}_2, ...\mathbf{S}_K$ such that every $\mathbf{x}_i$, $(i = 1, ..N)$ falls into one and only one of these regions.

$$\mathbf{S} = \mathbf{S}_1 \cup \mathbf{S}_2 \cup \mathbf{S}_3 .... \cup \mathbf{S}_K$$
$$\mathbf{S}_i \cap \mathbf{S}_j = \emptyset \qquad \forall\, i \neq j$$

6.6.1 Clustering Techniques

The clustering algorithms available in the literature can be broadly classified into one of two types [49]:

1. hierarchical or
2. partitional

Hierarchical clustering techniques proceed by either a series of successive mergers, or a series of successive divisions, imposing a hierarchical structure on the data. Agglomerative hierarchical methods start with the individual objects, thus there are initially as many clusters as objects. For Divisive hierarchical methods an initial single group of objects is divided until each object forms a group [37]. Spath [12] examines each of the Agglomerative and Divisive hierarchical clustering techniques in detail.

A *partitional clustering* technique organises the patterns into a small number of clusters by labeling each pattern in some way; these clustering techniques group items rather than variables into clusters. A pattern matrix is usually clustered in this way, which explains the extensive use of these techniques in pattern recognition and image processing [49]. Partitional techniques make use of criterion functions (square error), density estimators (mode seeking), graph structures and nearest neighbours [49]. Hierarchical clustering could in fact consist of a sequence of partitions. Detailed explanations, algorithm names and test results on both hierarchical and partitional clustering techniques can be found in [49].

Section 6.7: Rule Induction from Training Sets

A number of basically intuitive clustering algorithms, the details of which are often found in any chapter on Clustering in Pattern Recognition systems, are based on the Euclidean distance concept. They include

- Nearest-Neighbour algorithm
- Nearest-Centroid algorithm

Unlike the simple and basically intuitive algorithms above, two algorithms that are based on the minimisation of a performance index, (defined as the sum of the squared distances from all points in a cluster domain to the cluster centre) are the

- K-Means algorithm [39], [51], [37]
- Isodata algorithm [22]

The Clustering process is suitable when no *a priori* information or training sets for the relevant pattern classes exist. In the case of a dried fruit colour classification system, training sets that are manually selected by trained staff did exist. These *labeled sets* were necessary to specify the range of colours typical to each class or grade of dried fruit. For this reason the Clustering technique was not used.

6.7 Rule Induction from Training Sets

"Learning can be seen as a task of creating and maintaining an internal model of the world... One of the most constrained variants is the task of learning from a set of training examples, to predict a particular unknown property or 'class' of a situation from a set of known properties of that situation. This learning task is sometimes referred to as 'rule induction' from examples" [55]. In the most basic 'learning' system, the system assumes that each class is solely a function of the properties or attributes which the designer has used to describe the training examples. This is the approach that is used in this learning procedure: the training set 'teaches' the dried fruit classification system the colours upon which it is to base its decision rule.

Sample groups or *labeled sets* of dried fruit were hand selected by trained SADF [34] staff. For three fruit types: dried peach, dried apricot, and dried pear, 204, 204 and 128 sample units were selected respectively. The full data set or sample group of each fruit type was made up of four classes: Choice, Standard, Sub-Standard and Reject. In the case of dried peaches and dried apricots, there were 50 examples of each class; in the case of dried pears there were 32

examples of each class. In order to be able to test the final classifier, each *labeled set* was halved; one half of the labeled set (the *training set*) was used to train the classifier, which means the dried peaches in the *training set* were used to set up the parameters of the normal probability distributions for each class; the other half of the *labeled set* was used to evaluate the accuracy of this classifier. This method of classifier evaluation is called the **hold-half-out method** and is explained in the next section.

6.8 Evaluation of Class Separability and final Classifier method

6.8.1 Measuring the Separability of the classes :- Jeffries-Matusita distance

In order to assess the accuracy that we can expect from a particular pattern classification, it is desirable to have some measure of the resulting class separabilities. In other words, we need some indication of how much overlap there is between the pattern classes that were constructed from training sets. Although good class separability is not the only indication of well defined groups, a badly separated class is indicative of a bad choice of class defining variables. One such separability measure is the Jeffries-Matusita (J-M) distance [51]. This method provides a measure of the distance between any two classes in a given 2-dimensional coordinate system. Assuming that each class can be modelled by a multivariate normal distribution, the J-M distance between class i and class j is defined as :

$$\begin{aligned} \mathcal{F}_{ij} &= \sqrt{2(1-e^{-a})} \\ \text{where } a &= \frac{1}{8}(\mu_i-\mu_j)^T\left(\frac{\mathbf{\Sigma}_i+\mathbf{\Sigma}_j}{2}\right)^{-1}(\mu_i-\mu_j) + \frac{1}{2}(b) \\ b &= \log_e\left(\frac{\left|\frac{\mathbf{\Sigma}_i+\mathbf{\Sigma}_j}{2}\right|}{\sqrt{|\mathbf{\Sigma}_i|\,|\mathbf{\Sigma}_j|}}\right) \end{aligned} \tag{6.12}$$

and

μ_i is the mean value of all elements in class i
$\mathbf{\Sigma}_i$ is the covariance matrix of the data set or class i

$\mathcal{F}_{ij}$ gives a measure between two classes, and can be evaluated for each i, j where i and j range over all classes. In the case of a dried fruit having 4 classes, both i and j would range from 1 to 4. Classes that overlap with one another will have **low** J-M distance measures and consequently low accuracy can be expected from a classifier which is based on such class distributions. The J-M distance has an upper bound of $\sqrt{2}$, therefore a J-M distance measure of approximately 1.41 would indicate two perfectly separated classes of data.

Section 6.9: Classification of Dried Fruit data sets

The J-M distance measures were calculated for each pair of classes in each fruit type, and used as a pre-test of the success of the class separability where R (Red value) and G (Green value) are the class defining variables. These results are listed in Chapter 7.

6.8.2 Final Classifier Evaluation

The method selected to evaluate the classifier depends upon the availability of a *labeled set*. If a labeled set is available two similar approaches are available: the **hold-half-out method** [49], [38] and the **hold-one-out method**. In the *hold-half-out method* the total labeled set is divided in half; one half is used to construct the decision rule, the other half to test this decision rule. The roles of each set are then switched and the test repeated. Once each half has played the role of both 'training set' and 'test set', the error rates for both 'test sets' are combined to establish the final error estimate. This method is presumed to exaggerate the error rate of the classifier [45]. The *hold-one-out method* is similar to the *hold-half-out method*, except that on each occasion, only one element of the labeled set is selected as the test set, and the remainder of the labeled set form the training set. This method has a less pessimistic error rate prediction than the *hold-half-out method*, but is more processing intensive. Goncalves [7] discusses a similar 'Cross Validation' classifier evaluation method where blocks of samples rather than single items are left out of the test set. In this investigation, the *hold-half-out method* was adopted to evaluate the dried fruit classifier, the results of which constitute Chapter 7.

6.9 Classification of Dried Fruit data sets

The method of defining a dried fruit class using a normal probability distribution, was adopted. For each of the M classes of a particular type of dried fruit, a normal probability distribution would be defined. Labeled sets selected by trained SADF [34] staff, were used to identify the parameters of each probability distribution. For each fruit type M normal distributions were defined. The Jeffries-Matusita distance measure identifies the separability of each pair of classes. The Bayesian Classification method could have been used if *a priori* probabilities of class occurrence were known. At this point it should be noted that the Bayesian Classification method and the Quadratic discriminant method are almost identical. The two methods differ only in that the *a priors* of the Quadratic discriminant method are set equal in the absence of specific knowledge. Knowledge of these class probabilities for dried fruit (i.e. there will be a 60% chance of any dried fruit being in Choice grade, and therefore a 40% chance of a dried fruit being in Reject grade in a 2-class system), is however, unlikely. For this reason the Quadratic Discriminant was the most feasible option. The Quadratic

Discrimination Classification Rule of Equation 6.10, after careful investigation of the data sets, seemed to be the most suitable classification method for the task.

Section 6.9: Classification of Dried Fruit data sets

Chapter 7

Results

7.1 Introduction

This chapter is divided into three main sections; one for dried apricots, one for dried peaches, and one for dried pears. Each section deals only with the test results and conclusions for that particular dried fruit type. This approach was necessary because of the uniqueness of the processing requirements for each fruit type, as well as making the discussion of the results easier. The procedure used for all of the results in this chapter is the *Quadratic Discrimination Classification Rule* of Equation 6.10. Each manually selected control-group of dried fruit was used to both train and evaluate the Quadratic Discriminant Classifier (hereafter referred to as 'the Classifier') according to the *half hold-out* classifier evaluation method of Section 6.8.

Each dried fruit type (dried apricot, dried peach, dried pear) is divided into **four classes: Choice, Standard, Sub-Standard and Reject**. These classes are based entirely on the grading system used for dried fruit at South African Dried Fruit (SADF) [34]. Trained SADF staff manually selected examples of each dried fruit type. For dried apricots and peaches, 51 samples from each of the four classes were selected. For example, the full *data set* of dried apricots included four *labeled sets* :

- 51 Choice Apricots
- 51 Standard Apricots
- 51 Sub-Standard Apricots
- 51 Reject Apricots

In the case of dried pears, the *labeled sets* each consisted of 32 pears.

Section 7.1: Introduction

Half of each *labeled set* (i.e. 25 peaches, apricots and 17 pears) were used in the *training run* to set up the normal probability distribution for that particular dried fruit class. This was done for each of the four classes and resulted in a system of overlapping normal probability distributions for each fruit type, as described graphically in Figure 6.7 for a 3-class system. Each fruit in the remaining half of each *labeled set*, referred to as the *test set* was classified in order to gauge the success rate of the Classifier System.

Consider, for example, that apricot number 30 of the Choice class *labeled set* has a feature vector value of $\mathbf{x} = (\mathbf{200}, \mathbf{50})$. Upon substituting this vector into the four specified probability distributions, the following probability values were produced :

$$p(Choice) = 0.8$$

$$p(Standard) = 0.3$$

$$p(Sub - Standard) = 0.02$$

$$p(Reject) = 0.001$$

This particular apricot would be automatically classified as a Choice grade apricot according to the Quadratic Discrimination Classification Rule. As required by the **hold-half-out method**, the two sets were then swapped: the *training set* became the *test set*. This method results in an effective *test set* which is the size of the *labeled set*.

There are two 'RUNS' or parts to each section. RUN A tests were performed only upon those fruit classes or grades that were specified according to colour only. Those classes that included fruit which were graded according to **geometry** or **blemish size** were not included in RUN A tests. More specifically this meant that for dried apricots and peaches, RUN A tests consisted of three classes: Choice, Standard and Reject, and for dried pears RUN A tests consisted of two classes: Choice and Standard. RUN A was meant only to examine the success of the colour classification scheme. RUN B tests, on the other hand, were performed on all fruit classes. RUN B could be called a 'blind run' as it included all possible types of dried fruit; some, for example, were too small, some squashed, and some blemished. RUN B highlighted whether or not the grading of fruit could be done according to colour only. The same *training sets* were used in each of the runs i.e. the Choice *training set* used in RUN A is exactly the same as the Choice *training set* used in RUN B.

There is no general set of class selection criterion, or rather grade specifications, that hold for all dried fruit types. At South African Dried Fruit, each dried fruit type has its own particular set of rules for class allocation. Although the colour of the fruit is often relevant (sometimes

the only deciding factor), there are other factors that are important in class allocation. The different dried fruit class specifications that are listed in this chapter are those specified verbally by South African Dried Fruit [34]. Because colour is the major factor in dried fruit class allocation, the Classifier was designed for a colour feature space. The following results indicate the success of this method, but also indicate the necessity for a classification system based on more factors than colour alone.

The remainder of this chapter will be divided by fruit type. The results for each fruit type will be documented in the following way :

1. Manual Selection Criterion for each class
2. Colour plate examples of each dried fruit class
3. Jeffries-Matusita class separability measures between each of the four classes
4. COLOUR CLASSIFIER RESULTS - RUN A : Test results for those classes graded on colour only
5. COLOUR CLASSIFIER RESULTS - RUN B : Test results for all classes graded on colour, geometry, and blemish size
6. Conclusion

7.2 Dried Apricots

For each dried apricot a feature vector was set up to represent that particular apricot. This vector was 2-dimensional and consisted of the R value - the average red band value calculated over every pixel in the fruit, and the G value - the average green band value calculated over every pixel in the fruit. The images of the apricot were of either the front or the back (flat surface) of the fruit. Approximately half of the selected set are images of the front of the dried fruit, the other half are of the back of the fried fruit.

7.2.1 Manual Selection Criterion for each class - Dried Apricots

Choice grade Apricots

1. Bright orange in colour, with no variation in tone or colour.
2. No blemishes of any kind e.g. damage by hail, dark spots.

Section 7.2: Dried Apricots

3. No holes or other geometrical irregularities in apricot e.g. squashed fruit, ragged edges, too small in area.

Standard grade Apricots

1. Darker colour shade of orange than in Choice grade.
2. No blemishes of any kind e.g. damage by hail, dark spots.
3. No holes or other geometrical irregularities in apricot e.g. squashed fruit, ragged edges, too small in area.

Sub-Standard grade Apricots

1. No black blemishes.
2. Includes all apricots with holes or other geometrical irregularities in apricot e.g. squashed fruit, ragged edges, too small in area.
3. Some apricots are a darker shade of orange than those of Standard grade.

Reject grade Apricots

1. Any apricot that has a black blemish that is clearly visible to the eye.

COLOUR SAMPLES OF DRIED APRICOTS

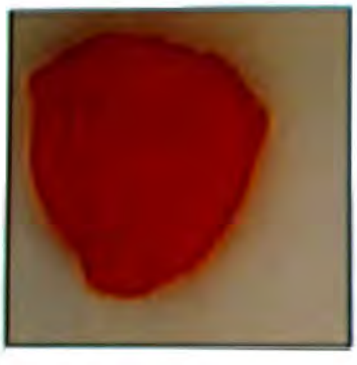

Choice grade

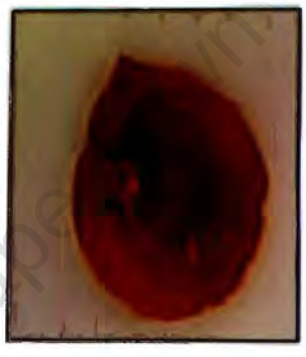

Standard grade

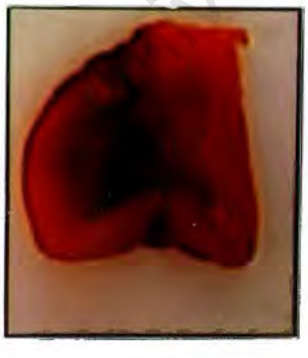

Sub-Standard grade

Reject grade

7.2.2 Jeffries-Matusita class separability measures between each of the 4 apricot classes

Jeffries-Matusita measures are a measure of the separability of two classes or cluster groups of data. Because each dried fruit class is represented by a normal probability distribution, the parameters which are used to calculate the J-M separability, are the class mean and covariance matrix. The maximum class separability is $\sqrt{2}$ which is approximately 1.41. A J-M separability measure of 1.41 indicates that two classes are completely separated, with no data overlap at all. Niblack [51] does not give any indication or examples of J-M separability measures that constitute 'poor' separability; the results of this chapter, however, indicate that clusters with J-M measures below the value of J-M = 1 are not well separated. A J-M separability measure is set up for two cases : the first is for the first *training set* of size 25; the second is for a *training set* of size 51, which is the entire *labeled set*. A *training set* much larger than 25 is recommended for implementation and any difference in J-M separability measure between the two cases should indicate the improvement (or lack thereof) in class separation that could be expected for larger *training sets*.

J-M Separability measures between classes of Dried Apricots						
Apricots	Choice Grade	%	Standard Grade	%	Sub-Std Grade	%
Std grade	1.307	93				
	(1.321)	*(94)*				
Sub-Std grade	0.99	70	0.53	38		
	(1.026)	*(73)*	*(0.542)*	*(38)*		
Reject grade	1.404	100	1.305	93	1.271	90
	(1.408)	*(100)*	*(1.346)*	*(95)*	*(1.267)*	*(90)*

The above table shows the Jeffries-Matusita separability measures between each of the dried apricot classes. The separability measures in brackets indicate the results from a *training set* of 50 apricots. The separability of the two classes is also indicated as a percentage, where a value of 1.41 would measure 100% separability. In most cases the class separability measure has improved for larger test *training sets*. The above J-M results will be discussed further in the 'RESULTS' section.

7.2.3 COLOUR CLASSIFIER RESULTS - DRIED APRICOTS

R U N A:

Dried fruit type : APRICOTS

Possible classes :

- CHOICE (C)
- STANDARD (S)
- REJECT/BAD (B)

A *labeled set* of 51 manually selected apricots exists for each of the three dried apricot classes. 25 of each of the *labeled set* were used in a *training run* to construct the normal probability distributions for each class. The remaining *test set* of 26 apricots were used to test the success rate of this 3-class–Colour-Classifier for dried apricots using the *Quadratic Discrimination Classification Rule* of Equation 6.10. The members of the *training set* and *test set* were then swapped, according to the **hold-half-out method** and the above procedure repeated. This process resulted in 51 dried apricots in each class *test set*, therefore a total of 204 apricots in the full test.

Dried Apricot Colour Classes

The feature vectors for each of the apricots in the 3 *labeled sets* for the classes of Choice, Standard, and Reject have been plotted in the (R, G) feature space in Figure 7.1. The ellipses are present merely to indicate the grouping of the vectors in each class. The clear separation between the three dried apricot classes within (R, G) feature space is evident in Figure 7.1.

Tables of Results for a 3 Class Colour Classifier System

The results have been divided into the three different classes of Choice(C), Standard(S), and Reject(B). In each case the number of apricots in the manually selected *test set*, is specified under 'Manually Allocated'. For example 51 apricots make up the *test set* for the Choice class. The classes to which the Colour Classifier allocated this same set of fruit is listed under 'Classifier Allocation'; the corresponding percentage of the total *test set* is listed under '% of Test Group'.

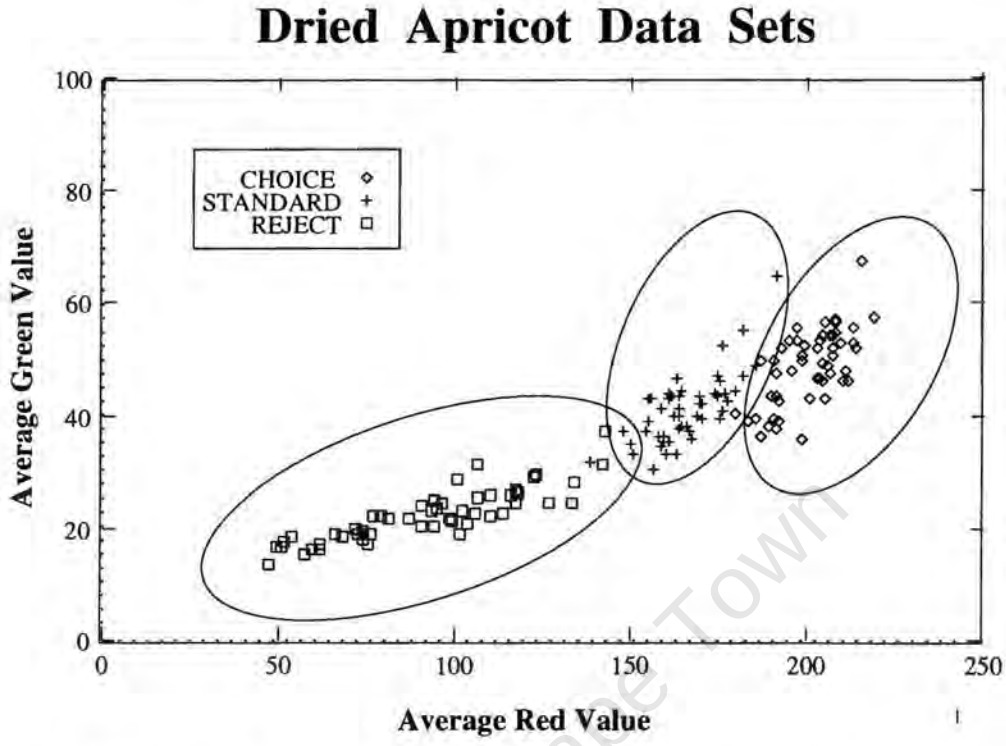

Figure 7.1: Feature vector distribution of 3 dried apricot classes

Allocation of Choice grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	C	49	C	96%
		2	S	4%
		0	B	0%

Allocation of Standard grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	S	50	S	98%
		1	C	2%
		0	B	0%

Section 7.2: Dried Apricots

Allocation of Reject grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	B	49	B	96%
		0	C	0%
		2	S	4%

7.2.4 Discussion of Results from RUN A

The Colour Classifier succeeded extremely well in distinguishing between the Choice, Standard, and Reject classes of dried apricots. The excellent separation between the 3 classes is confirmed by the J-M class separability measures for those classes: 93% separability between Choice and Standard class data(or clusters) ; 93% separability between Standard and Reject class data ; 100% separability between Choice and Reject class data.

All of the J-M class separability measures improve when each of the *training sets* is increased from 25 to 51 apricots. An even larger *training set* may achieve even better class separabilities. Because these *training sets* are used to define the parameters of the normal distribution for each class, the larger the *training set* set used, the more accurate the final distribution, and hence the more accurate the final Classifier system.

In the following run, RUN B, a Sub-Standard class of data is included. The normal distributions for four classes of data will now be calculated resulting in a 4-class-Colour-Classifier. This means that any new dried apricot (i.e. any apricot which is not part of the given *training set*) could be allocated to one of four classes.

R U N B:

Dried fruit type : APRICOTS

Possible classes :

- CHOICE (C)
- STANDARD (S)
- SUB-STANDARD (SS)

- REJECT/BAD (B)

A *labeled set* of 51 manually selected apricots exists for each of the four dried apricot classes. The Choice, Standard and Reject data sets used in RUN A were used again in RUN B. The Sub-Standard class *labeled set* is included in this test run.

Dried Apricot Colour Classes

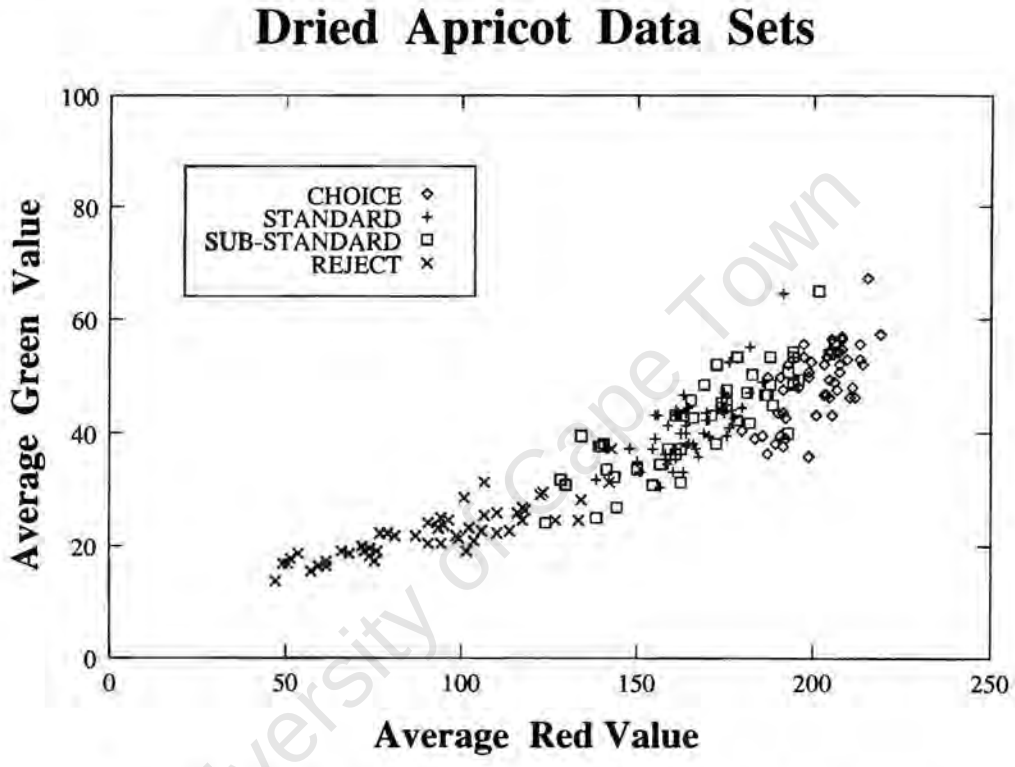

Figure 7.2: Feature vector distribution of 4 dried apricot classes

The feature vectors for each of the 51 apricots in each of the four *labeled sets* (Choice, Std, Sub-Std, Reject) have been plotted in the (R, G) feature space in Figure 7.2.

It is evident from Figure 7.2 that the Sub-Standard data encroaches rather hap-hazardly into all three of the other classes. This can be expected as Sub-Standard apricots are largely categorised according to their geometrical flaws rather than the flesh colour: any apricot with geometric flaws will always fall into the Sub-Std class. An apricot with a hole, for example, in a manual system would be allocated to the Sub-Standard class, regardless of its colour. The colour of this same apricot, if analysed, could in fact be that of a Choice apricot. (The characteristics which describe a Sub-Standard apricot are listed in Section 7.2.1).

Tables of Results for a 4 Class Colour Classifier System

Section 7.2: Dried Apricots

The results have been divided into the four different classes. As for RUN A, both the Manually designated class and the Classifier allocations are shown.

Allocation of Choice grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	C	42	C	82%
		1	S	2%
		8	SS	16%
		0	B	0%

Allocation of Standard grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	S	6	S	12%
		0	C	0%
		45	SS	88%
		0	B	0%

Allocation of Sub-Standard grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	SS	40	SS	78%
		7	C	14%
		2	S	4%
		2	B	4%

Allocation of Reject grade Dried Apricots				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	B	51	B	100%
		0	C	0%
		0	S	0%
		0	SS	0%

7.2.5 Discussion of Results from RUN B

As indicated by the J-M class separability measures, the Sub-Standard class was not well separated from the other classes: a J-M measure of 90% was achieved between the Sub-Standard (SS) and Reject class data ; a separability of 70% was achieved between the Sub-Standard and Choice data ; a J-M separability of only 38% was achieved between the Standard and Sub-Standard class data. By increasing the class *training sets* from 25 to 51 apricots, an increase in separability occurred between the Sub-Standard and Choice classes; the other class separabilities remained constant.

In the scatter-plot of Figure 7.2 it is clear that the data of the SS colour class is scattered throughout the other three classes. The normal distribution associated with this data set spread well into the other three classes, resulting in the confusion of apricot allocation seen in the results of RUN B.

The Sub-Standard class of Dried apricots includes mainly geometrically faulty apricots; they can have a range of colours - from Choice grade through to Reject grade. Each new *training set* of Sub-Standard apricots could be completely unlike any of the other *training sets*. In other words, because the colour of the fruit is not always considered, a large variety, and therefore range of apricot colours will exist in each SS set. The (R, G) colour space of the Sub-Standard class cannot, therefore, be described as a Normal distribution. This implies that all apricots allocated according to geometry should not be part of the SS *training set*.

7.2.6 Conclusion - Colour Classification of Dried Apricots

When dried apricots were divided into the three classes of Choice, Standard and Reject, (as specified by the control-groups selected by S. A. Dried Fruit) the Colour Classifier allocated apricots to the correct class at least 96% of the time. These results were achieved using a *training set* of only 25 apricots. *Training sets* of 100 or more apricots should lead to better defined clusters and consequently improved apricot allocation.

Handling of the Sub-Standard class :

The most obvious geometrically unacceptable apricots were removed from the initial Sub-Standard class *labeled set*. Upon substituting this set of data into the 3-class-Colour-Classifier, it was found that most SS apricots were allocated to the Standard class. This would indicate two options for processing of Sub-Standard apricots :

1. Apricot geometry must be tested first, either manually, or automatically, in which case relevant software procedures would have to be developed to handle this. There would be 3 apricot colour classes :- that of Choice grade, Standard grade and Reject grade. The apricots allocated to the Standard class should then be re-sorted manually and a group of Sub-Standard fruit derived from this.
2. A Sub-Standard colour class set up so that there would be four grades (or classes). The *training set* for this group should contain only geometrically acceptable apricots that describe the unique colour of the Sub-Standard class. An initial run to remove all geometrically flawed apricots would still be required prior to this colour classification procedure.

7.3 Dried Peaches

For each dried peach an (R, G) feature vector was set up to represent a particular peach. Unlike the apricot, not all of the fruit image pixels were used to find these average R, G values. Because the peach core area is similar in colour to a rotten peach colour, the core area was removed from the peach image (see Chapter 5), thereafter the remaining pixels in the peach image were averaged to get the R and G values for that peach. The images of the peach were of either the front or the back (flat surface) of the fruit. Approximately half of the images that produced the *labeled set* were images of the front of a dried peach, the other half were of the back of a dried peach.

7.3.1 Manual Selection Criterion for each class - Dried Peaches

Choice grade Peaches

1. Orange(darker shade than choice apricot) in colour, with slightly darker coloured edge or border.
2. No black/dark spots of any kind.

COLOUR SAMPLES OF DRIED PEACHES

Choice grade

Standard grade

Sub-Standard grade

Reject grade

3. No holes or other geometrical irregularities in peach e.g. flesh not thick enough, cut edges.

Standard grade Peaches

1. Darker, weaker colour shade of orange than in Choice grade.
2. No blemishes of any kind e.g. damage by hail, dark spots.
3. Any peach that has Choice coloured flesh, but has a geometrical irregularity i.e cut edge.

Sub-Standard grade Peaches

1. Darker, weaker colour shade of orange than in Standard grade.
2. Includes all peaches with holes or other geometrical irregularities e.g ragged edges, marks from hail or insects.
3. Any peach that has Standard coloured flesh, but has a geometrical irregularity i.e cut edge.

Reject grade Peaches

1. Darker colour shade of orange than in Sub-Standard grade.
2. Any peach that has a black blemish that is clearly visible to the eye.
3. Includes all peaches with stones, pips, or other objects that were not removed from the flesh by washing.

7.3.2 Jeffries-Matusita class separability measures between each of the 4 Dried Peach classes

A J-M separability measure is set up for two cases : the first is for a *training set* of size 25; the second is for a *training set* of size 50. A *training set* much larger than 25 is recommended for implementation.

J-M Separability measures between classes of Dried Peaches						
Peaches	Choice Grade	%	Standard Grade	%	Sub-Std Grade	%
Std grade	1.03	73				
	(1.024)	*(73)*				
Sub-Std grade	1.318	94	0.956	68		
	(1.286)	*(91)*	*(0.902)*	*(64)*		
Reject grade	1.34	95	1.156	82	0.535	38
	(1.355)	*(96)*	*(1.156)*	*(84)*	*(0.658)*	*(47)*

The above table shows the Jeffries-Matusita separability measures between each of the Dried Peach classes. The separability measures in brackets indicate the results from a *training set* of 50 apricots. The J-M measure between the Sub-Standard and Reject class is the worst at 38%. A larger *training set* of 50 peaches resulted in a class separability measure of 47%. One would expect a further improvement for a much larger *training set.*

The above results will be discussed further in the 'RESULTS' section.

7.3.3 COLOUR CLASSIFIER RESULTS - DRIED PEACHES

R U N A:

Dried fruit type : PEACHES

Possible classes :

- CHOICE (C)
- STANDARD (S)

- REJECT/BAD (B)

A *labeled set* of 51 manually selected peaches exists for each of the three dried peach classes.

Dried Peach Colour Classes

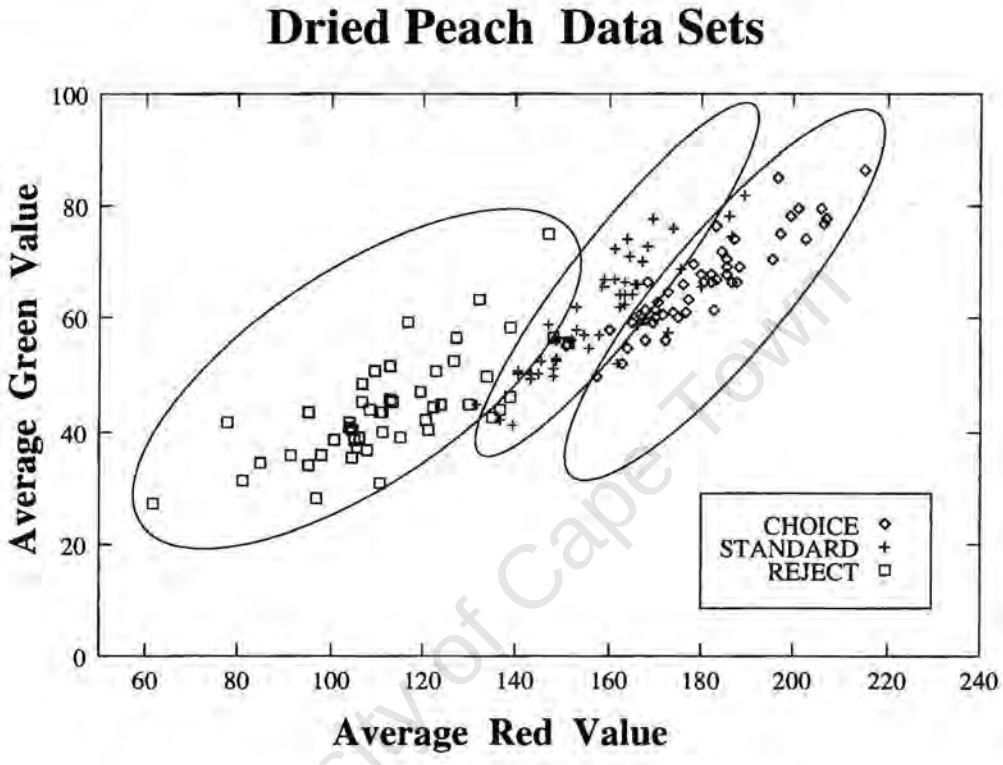

Figure 7.3: Feature vector distribution of 3 dried peach classes

The feature vectors for each of the 51 peaches in the three *training sets* (Choice, Std, Reject) have been plotted in the (R, G) feature space in Figure 7.3. The ellipses surrounding certain data are included only to emphasise the grouping of feature vectors within each class. Good separation between the three dried peach classes in (R, G) feature space is evident.

Tables of Results for a 3 Class Colour Classifier System

The results have been divided into the three different classes of Choice(C), Standard(S) and Reject(B). For each class both the Manually designated class and the Classifier allocations are shown.

Allocation of Choice grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	C	44	C	86%
		7	S	14%
		0	B	0%

Allocation of Standard grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	S	43	S	84%
		6	C	12%
		2	B	4%

Allocation of Reject grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	B	47	B	92%
		0	C	0%
		4	S	8%

7.3.4 Discussion of Results from RUN A

The Classifier worked well in distinguishing between Choice, Standard and Reject classes of peaches. The J-M class separability measure between the Choice and Standard class was only 73%, yet 86% of Choice peaches were correctly allocated, and 84% of the Standard peaches were correctly allocated. Although the Classifier allocated correctly at least 84% of the time, it must be remembered that these tests offer a pessimistic evaluation of the success of the Classifier. The J-M class separability measures were 82% between Standard and Reject classes, and 95% between the Choice and Reject classes. The very good separability between these classes is reflected in the test run results; 92% of the Reject peaches have been correctly allocated, but 8% of the Reject class were misclassified as Standard class, and no Reject peaches were allocated to the Choice class.

A reasonably small number of peaches were incorrectly allocated by the Classifier. If the *training sets* were much larger, a larger and more thorough class colour set should be defined, ensuring less misallocations.

R U N B:

Dried fruit type : PEACHES

Possible classes :

- CHOICE (C)
- STANDARD (S)
- SUB-STANDARD (SS)
- REJECT/BAD (B)

A *labeled set* of 51 manually selected peaches exists for each of the four dried peach classes. The Choice, Standard and Reject data sets used in RUN A were used again in RUN B.

Dried Peach Colour Classes

The feature vectors for each of the 51 peaches in each of the four *training sets* (Choice, Std, Sub-Std, Reject) have been plotted in the (R, G) feature space in Figure 7.4. As can be seen in the scatterplot of Figure 7.4, it is difficult to isolate four independent clusters of data. The data of the SS colour class is scattered throughout the other classes. This data distribution promised less successful results in peach allocations, than those found in RUN A.

Tables of Results for a 4 Class Colour Classifier System

The results have been divided into the four different classes of Choice(C), Standard(S), Sub-Standard(SS) and Reject(R). As for RUN A, both the Manually designated class and the Classifier allocations are shown.

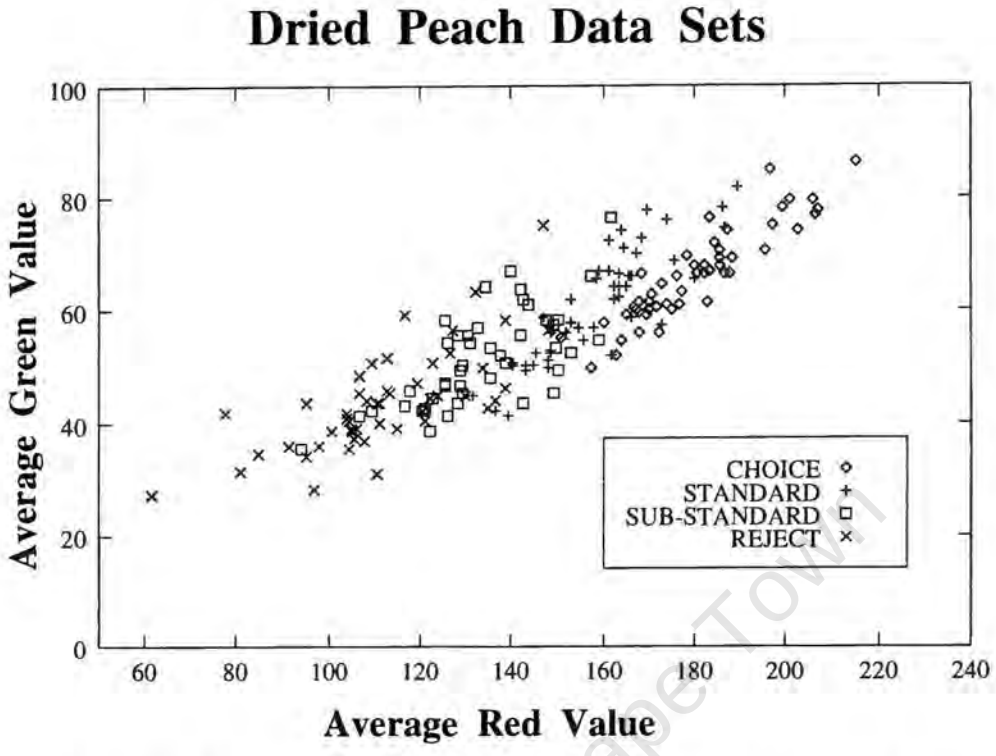

Figure 7.4: Feature vector distribution for 4 classes of dried peaches

Allocation of Choice grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	C	44	C	86%
		7	S	14%
		0	SS	0%
		0	B	0%

Allocation of Standard grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	S	37	S	72%
		6	C	12%
		8	SS	16%
		0	B	0%

Allocation of Sub-Standard grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	SS	22	SS	42%
		0	C	0%
		9	S	18%
		20	B	40%

Allocation of Reject grade Dried Peaches				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
51	B	41	B	80%
		0	C	0%
		2	S	4%
		8	SS	16%

7.3.5 Discussion on Results from RUN B

As indicated by the J-M class separability measures, the Sub-Standard class was extremely well separated from the Choice class with a measure of 94%; reasonably separated from the Standard class with a J-M measure of 68% and badly separated from the Reject class with a J-M measure of 38%. As expected from the J-M separability measures, the Classifier correctly allocated to the Choice class 86% of the time, as it did in the 3-class-Colour-Classifier case, but allocations in both Standard and Reject classes were less successful, due in both cases to peaches being allocated incorrectly to the Sub-Standard class.

Larger *training sets* improved all J-M class separabilities, except for those of the Sub-Standard class. This is understandable, as the fruit in the SS class ranges throughout the colour range of the C, S, and B classes, as seen in Figure 7.4, and larger training sets will merely increase the colour possibilities of the class.

7.3.6 Conclusion - Colour Classification of Dried Peaches

If dried peaches are graded into classes of Choice, Standard, and Reject, then the Colour Classifier was at least 84% successful when a training set of only 25 peaches was used to define the parameters of each normal distribution. As indicated by the improved J-M separability measures for a training set of 51 dried peaches, a larger training set would probably increase the success rate of the colour classification scheme.

Because of the geometrical irregularities present in many of the dried peaches allocated to the SS class, the Classifier which classifies according to colour alone, is not adequate. An initial run which would test the geometry of each of the dried peaches, is necessary. Only geometrically acceptable dried peaches would be processed and graded by the Classifier; the remainder would be graded as Sub-Standard dried peaches. Because the Sub-Standard dried peaches consist almost entirely of geometrically flawed fruit, it is probable that a Sub-Standard colour class, as in the case of dried apricots, would be unnecessary.

7.4 Dried Pears

As in the case of the dried apricots all of the pear image pixels were used to find the average values (R, G) of each fruit. Again, images of both the front and back of the dried pears were used. The following results are those achieved using the **hold-half-out** classifier evaluation method, with an effective data test set of 32 dried pears, in each class, or a full test set of 128 dried pears.

7.4.1 Manual Selection Criterion for each class - Dried Pears

Choice grade Pears

1. A specific yellow/green in colour, with no variations in tone or colour.
2. No blemishes of any kind e.g. damage by hail, dark spots.
3. No holes or other geometrical irregularities in apricot e.g. squashed fruit, ragged edges, too small in area.

Standard grade Pears

1. Darker colour shade of yellow/green than in Choice grade.
2. No blemishes of any kind e.g. damage by hail, dark spots.

COLOUR SAMPLES OF DRIED PEARS

Choice grade

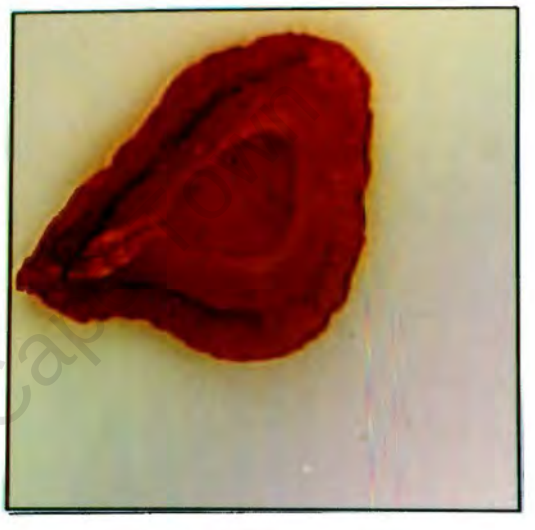

Standard grade

Sub-Standard grade

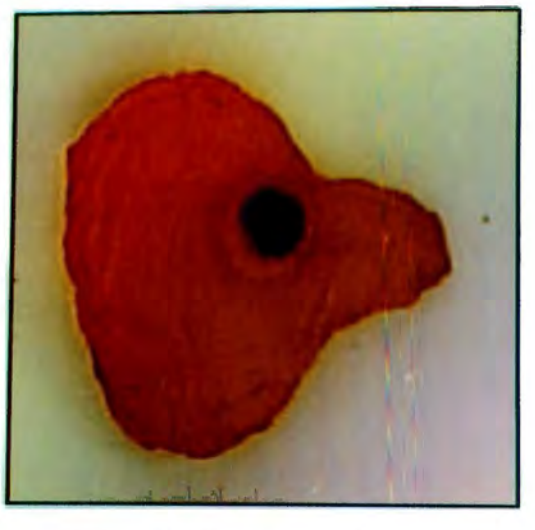

Reject grade

3. No holes or other geometrical irregularities in apricot e.g. squashed fruit, ragged edges, too small in area.

Sub-Standard grade Pears

1. Includes all pears with holes or other geometrical irregularities e.g. squashed fruit, ragged edges, too small in area.
2. Some apricots are a darker shade of orange than those of Standard grade.
3. Any pear that has one or more blemishes of any kind. Black blemishes must have a diameter of less than a given value e.g. 6 mm.
4. Any pear that was not fully ripe when picked, and has a less 'full ripe' colour to the trained eye.

Reject grade Pears

1. Any pear that has one or more black blemishes that are greater than a given diameter (e.g. 6 mm).

7.4.2 Jeffries-Matusita class separability measures between each of the 4 pear classes

A J-M separability measure is set up for two cases : the first is for a *training set* of size 16; the second is for a *training set* of size 32. A *training set* larger than 100 is recommended for implementation.

J-M Separability measures between classes of Dried Pears						
Pears	Choice Grade	%	Standard Grade	%	Sub-Std Grade	%
Std grade	1.119	79				
	(1.125)	*(80)*				
Sub-Std grade	0.61	43	0.667	43		
	(0.531)	*(38)*	*(0.706)*	*(50)*		
Reject grade	0.828	59	0.71	50	0.383	27
	(0.802)	*(57)*	*(0.72)*	*(52)*	*(0.330)*	*23)*

The above table shows the Jeffries-Matusita separability measures between each of the dried Pear classes. The separability measures in brackets indicate the results from a *training set* of 32 pears. The above J-M separability measures indicate that only the Choice and Standard classes have a reasonable measure of separability.

7.4.3 COLOUR CLASSIFIER RESULTS - DRIED PEARS

R U N A:

Dried fruit type : PEARS

Possible classes :

- CHOICE (C)
- STANDARD (S)

A *labeled set* of 32 manually selected pears exist for each of the two dried pear classes. 16 of this *labeled set* have been used in the *training run* to construct the class normal probability

distributions. The remaining pears have been used to test the success rate of this 2-class-Colour-Classifier.

Dried Pear Colour Classes

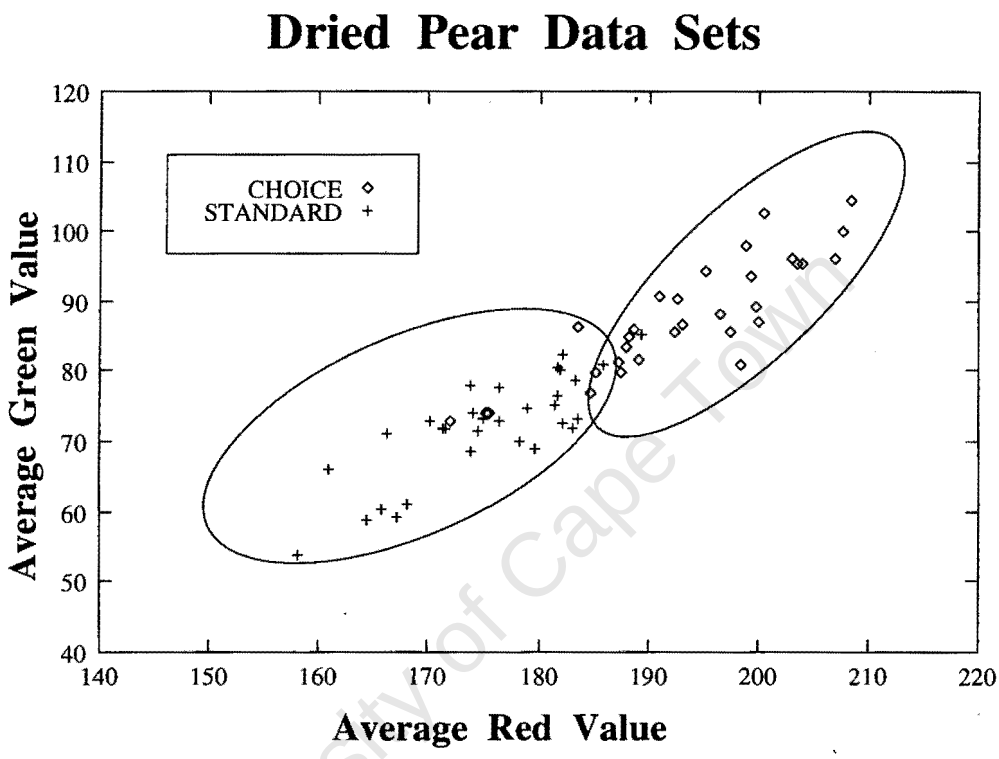

Figure 7.5: Feature vector distribution for 2 classes of dried pear

The feature vectors for each of the 32 pears in each of the two *training sets* (Choice, Standard) have been plotted in the (R, G) feature space in Figure 7.5.

Tables of Results for a 2 Class Classifier System

The results have been divided into the two different classes. For each class both the Manually designated and the Classifier allocations are shown.

Section 7.4: Dried Pears

Allocation of Choice grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	C	29	C	91%
		3	S	9%

Allocation of Standard grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	S	30	S	94%
		2	C	6%

7.4.4 Discussion on Results from RUN A

The only promising J-M class separability measure, of 79%, was between the Choice and Standard colour classes. A larger *training set* of 32 pears improved the separability to 80%. The J-M separability measures between all other pear classes are unsatisfactory. It is unlikely that larger *training sets* would have improved these pear class separabilities to the region of 90%.

In RUN A, where only the two pear classes of Choice and Standard exist; 91% of Choice pears were correctly classified, and 94% of Standard pears were correctly classified.

R U N B:

Dried fruit type : PEARS

Possible classes :

- CHOICE (C)
- STANDARD (S)
- SUB-STANDARD (SS)
- REJECT/BAD (B)

A *labeled set* of manually selected pears exist for each of the four dried pear classes. The Choice and Standard data sets used in RUN A are used again in RUN B.

Dried Pear Colour Classes

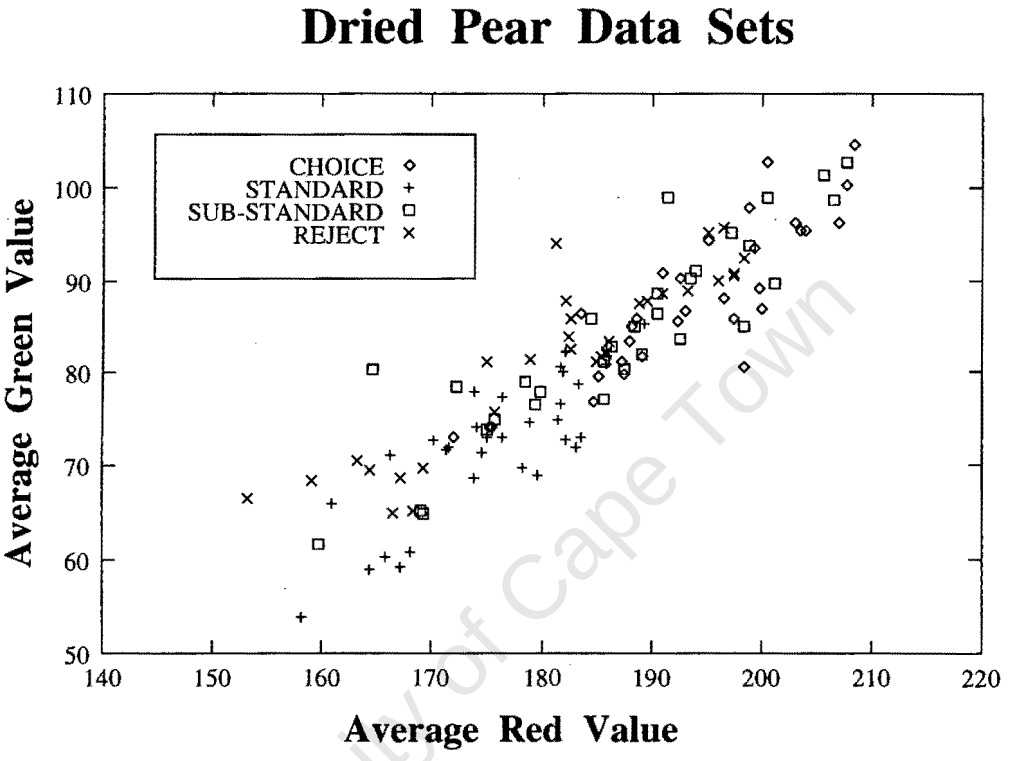

Figure 7.6: Feature vector distribution for 4 classes of dried pear

The feature vectors for each of the 32 peaches in each of the 4 *training sets* (Choice, Std, Sub-Std, Reject) have been plotted in the (R, G) feature space in Figure 7.6.

Tables of Results for a 4 Class Colour Classifier System

The results have been divided into the 4 different classes. As for RUN A, both the Manually designated and the Classifier allocations are shown.

Section 7.4: Dried Pears

Allocation of Choice grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	C	15	C	47%
		0	S	0%
		14	SS	44%
		3	B	9%

Allocation of Standard grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	S	17	S	53%
		0	C	0%
		5	SS	18%
		10	B	29%

Allocation of Sub-Standard grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	SS	13	SS	40%
		10	C	31%
		3	S	10%
		6	B	19%

Allocation of Reject grade Dried Pears				
Manually Allocated		Classifier Allocation		% of Test
Number	Class	Number	Class	Group
32	B	18	B	56%
		1	C	3%
		2	S	6%
		11	SS	34 %

7.4.5 Discussion on Results from RUN B

In RUN B where a Sub-Standard and Reject class are included, 56% or less of any one group of pears were allocated to the correct class.

7.4.6 Conclusion - Colour Classification of Dried Pears

The method of classifying pears according to the (R, G) values of the pear is successful for Choice and Standard grades of dried pear. Classification by colour alone is unsuccessful in the Sub-standard and Reject classes. Sub-Standard and Reject class pears are defined by the presence and size of one or more black blemishes. The remainder of the pear flesh, in these two classes, would often be described as Choice grade in colour. The width of the black blemish(es) on the pear flesh is used to divide the pears into Sub-Standard and Reject classes. To ensure successful classification of dried pears, the blemish areas (and width) must be processed and considered independently. A method of blemish isolation is discussed in Chapter 5. It was found that no geometry test for damaged pears was necessary; pears, once dried are a hardy fruit and not one geometrically faulty pear was found in the full test set of 128 dried pears.

Finally, the process of dried pear classification must consist of two areas :

1. blemish isolation and blemish sizing
2. colour classification of non-blemished pear flesh

The process of isolating all blemishes has been implemented and tested on the same training set of dried pears that were used in the colour classification system tests. The results of this process are detailed in the next section.

7.4.7 Colour and Blemish Analysis Classifier System

Choice and Standard grade pears are currently manually graded according to flesh colour only. Sub-Standard pears are graded mainly according to the presence of black blemishes, and to a very small extent, by their colouring. Reject pears are those with black blemishes larger than a given diameter. Because of the above grading methods, only three possible pear colour types exist; Choice, Standard and Sub-Standard. When evaluating the normal

distribution parameters for the SS colour class, clearly only those pears that have no black blemishes, must be selected for the *training set.*

In the proposed dried pear grading process, the first step of the process would be to isolate all blemished pears, thereby ensuring that only Choice, Standard and some Sub-Standard pears remained. The pears that 'passed' the blemish test would then be graded by colour alone.

Of the 32 pears that made up the manually selected group of SS class pears, 31 had black blemishes, and the remaining pear had dark orange colouring. The process of isolating those Sub-Standard pears with black blemishes, described in Chapter 4, Section 4.5, was entirely successful; the full blemish area of every single blemish on every single pear was identified. Similarly, all blemishes on the group of Reject class pears were successfully isolated. A number of Choice grade dried fruit were submitted for blemish testing, and no blemishes were found.

7.4.8 Conclusion - Colour and Blemish Classification of dried Pears

Full results for this process could not be done; a reasonably sized *training set* of Sub-Standard coloured pears was required to construct a reasonable normal distribution for that colour class:- this was not available. However, it is reasonable to predict the following :

- Using the blemish isolation method of Chapter 4, nearly all black blemishes should be successfully isolated, promising a high success rate in allocation of dried pears to Sub-Standard and Reject classes.
- Using only two colour classes, 91% of Choice grade dried pears were correctly allocated, and 94% of Standard grade dried pears were correctly allocated. This would imply that the proposed colour classifier should be very effective in grading these classes of dried pears by colour.

The independent results achieved using the Blemish Isolation procedure as well as the Colour Classifier are far better than those achieved using the Colour Classifier alone. In order that successful dried pear classification occurs, it is recommended that an initial blemish isolation and blemish sizing process is performed on each dried pear. Depending on the blemish size, the pear would be allocated to either the Sub-Standard or Reject class. Those dried pears that were found to be blemish-free, would be submitted to the Colour classification process for classification into the Choice, Standard or Sub-Standard classes. It was found however, that the process of blemish analysis was very sensitive - dried pears needed to be washed effectively before grading.

Chapter 8

Conclusion and Recommendations

8.1 Introduction

Numerous machine vision systems are available, which successfully grade fresh fruit; only one such machine has been developed in South Africa. The available machine vision systems are, however, only partially successful in grading dried fruit. The goal of this dissertation was the investigation of a pattern recognition system which would successfully analyse and classify dried fruit.

8.2 Conclusions

A human classification of a dried fruit into a particular grade is based upon the colour, shape, size, and blemish detail of the dried fruit. In order, therefore, for a machine vision dried fruit grading system to be entirely successful, colour image classification will have to be combined with blemish and shape analysis.

8.2.1 The Feature Space

A feature vector had to be found which would accurately and successfully relate any dried fruit to its respective grade. Because colour is a major consideration in the visual grading of dried fruit, it was decided that a colour system would be the best defining feature space for the problem. An investigation of literature and current machine vision fruit grading systems, highlighted the RGB and HSI colour systems. A comparison of these two colour models resulted in the selection of the Red(R) and Green(G) coordinates of the RGB colour space.

8.2.2 Discussion of Segmentation Issues

Two segmentation issues were imperative to the success of the dried fruit classification system. These were :

- The removal of the dark red/brown coloured core area from dried peaches.
- The isolation and sizing of Black blemishes upon dried pears.

The Morphological Top-Hat Transform process was found to be both the most efficient, and entirely successful method for removing the dark red/brown coloured core area from dried peaches. Basic automatic thresholding of the image colour histograms (or rather the Red and Green bands) was found to be very successful in isolating every single black blemish on each of the 42 dried pears that were tested.

8.2.3 Selection of a Pattern Classification Method

A pattern classification method which classified a sample fruit according to its feature vector variables Red(R) and Green(G), and a specific decision rule, had to be selected. Upon careful consideration of the class density distributions in the R, G feature space, it was found that the Normal density distribution best described the parameters of the class labeled sets. Because of the lack of **relevant** *a priori* probability distributions, the Bayes Classifier was not selected; the Quadratic Discriminant function was found to be the best classification method. Assumed *a priori* probabilities were used in this method. The same approach could in fact have been taken, using a Bayes Classifier:- the Bayes Quadratic Classifier is in fact very similar to the Quadratic Discriminant Classifier, as it assumes normal density distributions for the data classes, and has a Quadratic function form. The Bayes method is, however, best suited to cases where the actual *a priori* probabilities are known.

8.2.4 Results from the Quadratic Discriminant Classifier

The Quadratic Discrimination Classification Rule, based upon normal class data probability distributions, equal *a priori* probability ratios and equal misclassification cost ratios, was tested upon a total test set of 540 dried fruit using the pessimistic 'hold-half-out' Classifier Evaluation method. The results were very good for those classes of dried fruit which were based upon colour characteristics only. The blemish isolation and sizing procedure on dried pears was very successful; inclusion of this procedure in the final Pattern Recognition Classifier System promises to improve the overall allocation results substantially. It was found that the geometry of the dried fruit was very relevant in the Sub-Standard classes of dried apricot and

dried peaches, consequently an initial fruit geometry test would further improve the success of a Pattern Recognition Classification System for dried fruit.

The following simple flow-charts represent the recommended procedures necessary for a successful dried fruit classification system. For each dried fruit type that has been tested, there is a flow-chart describing the procedures that will produce the optimum results.

8.3 Recommendations

If a fully automated machine vision grading system for dried fruit is required, a method should be found to analyse the geometry of dried fruit. In this test set of dried apricots, dried peaches and dried pears, only the dried pears did not require a fruit geometry test. Because of the relatively round shape of dried apricots and dried peaches, a test of the *convexivity* of their borders may prove to be an ample test for a suitable geometry. This solution, however, is particular to dried fruit with 'rounded' shapes; it may not be suitable for another shape of dried fruit. The area of the fruit would also have to be tested in order to find whether a dried fruit was too small, or too large. Further investigations would be required in both of these areas.

If this Pattern Recognition Colour Classification system was implemented, much larger training sets of dried fruit should be used in the training run to define the normal class density parameters. This would ensure that many of the variations of dried fruit colour within each grade would be catered for in the training run.

The Quadratic Classification Rule of the Pattern Recognition Colour Classification system could be more successful and relevant if :

1. specific *a priori* class probability ratios were known, and not assumed to be equal
2. relevant misclassification cost ratios were set up, and not assumed to be equal.

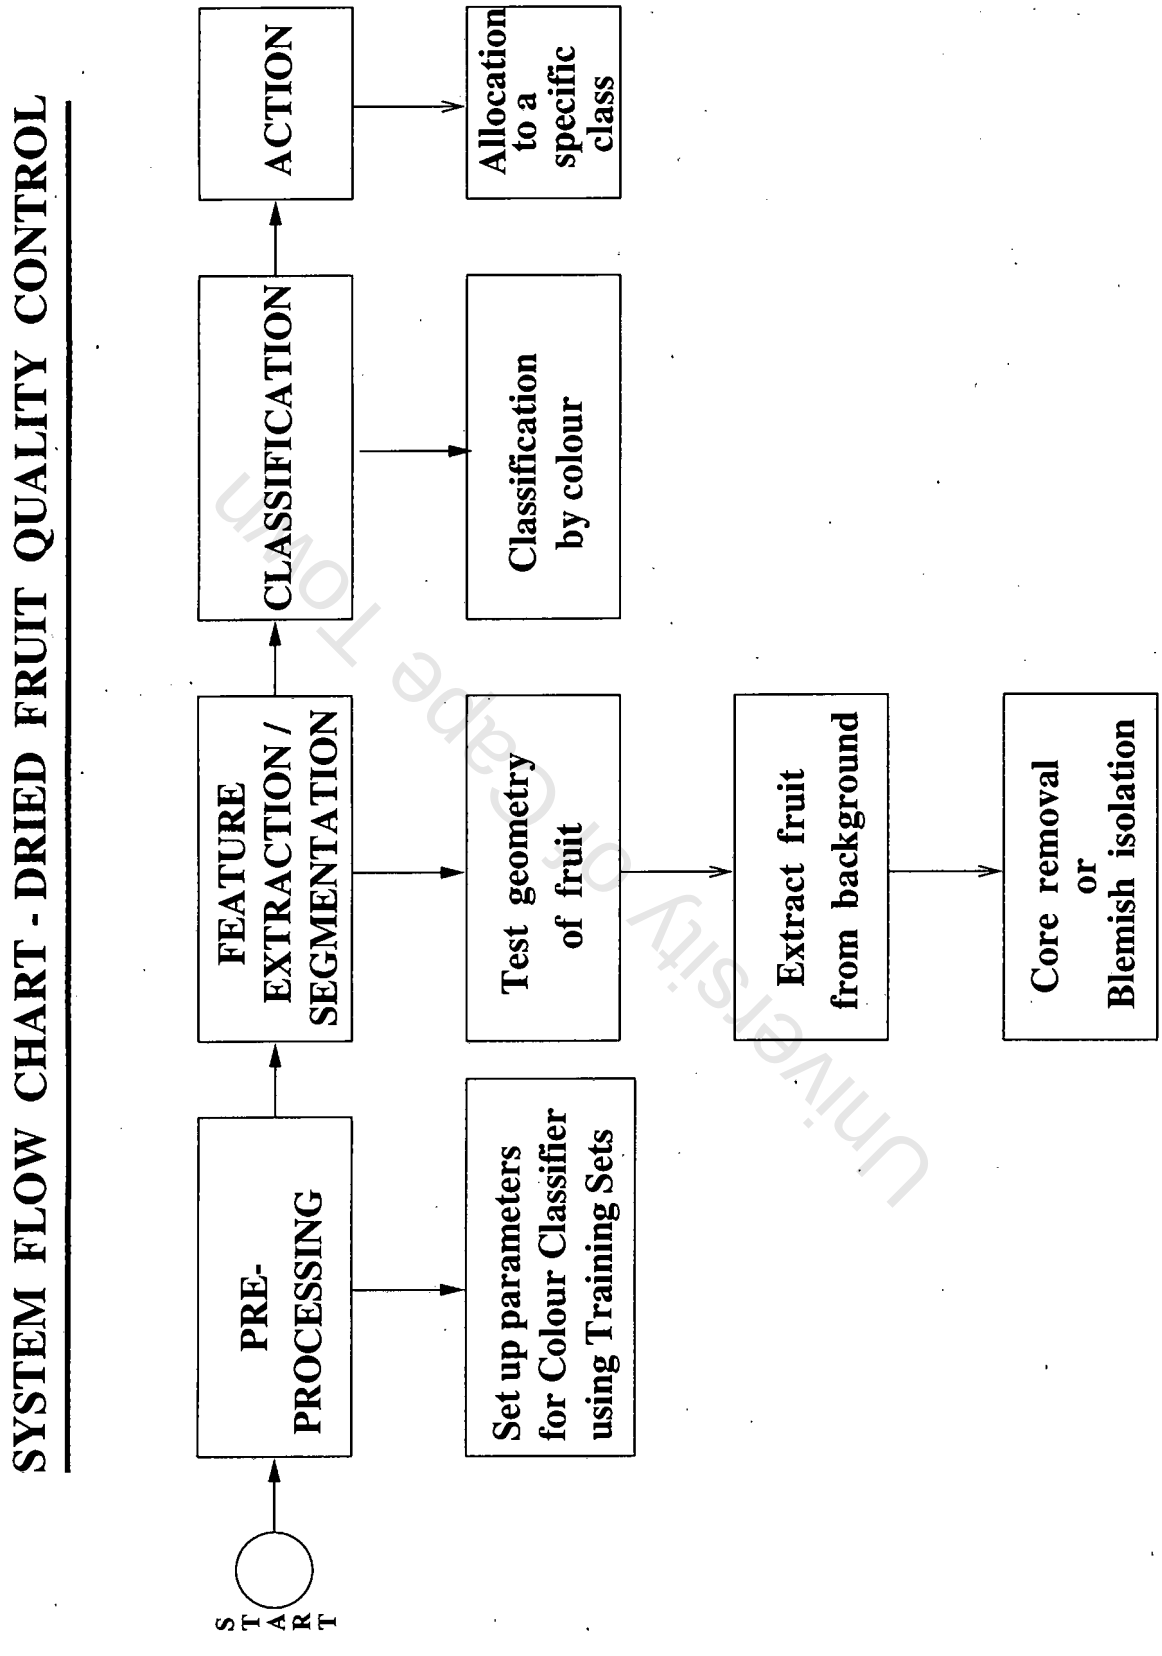

Figure 8.1: System Flowchart for an automated Dried Fruit Classification System

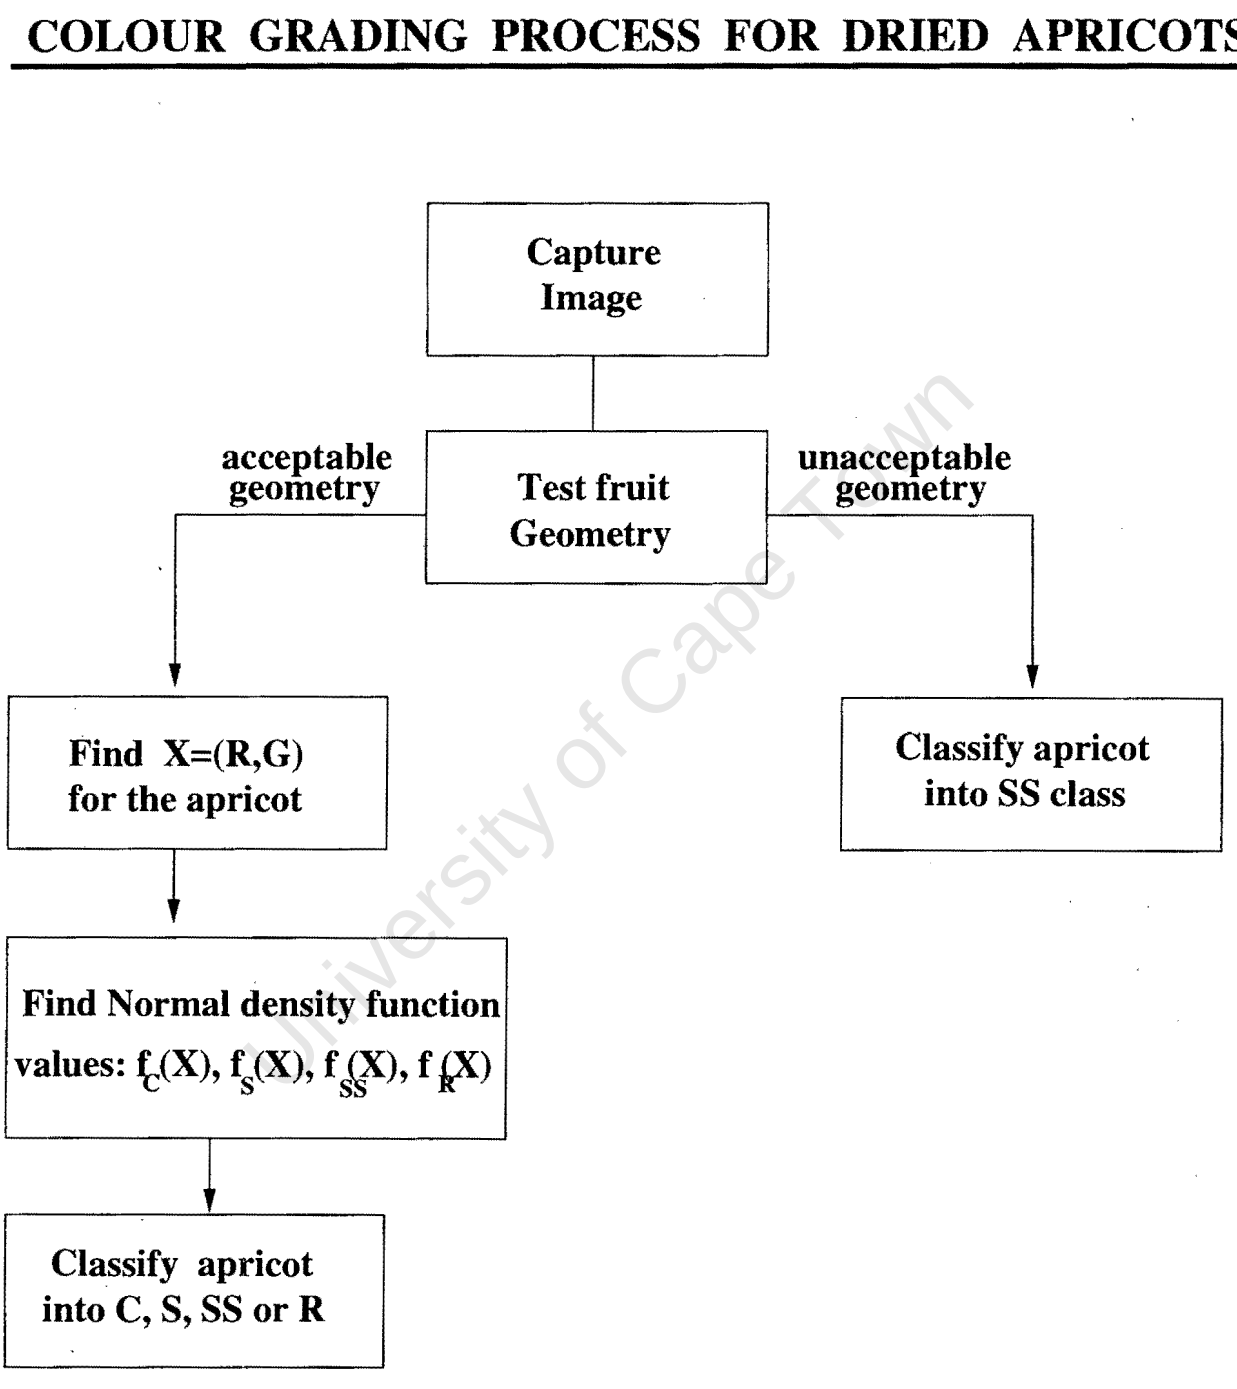

Figure 8.2: System Flowchart for classification of Dried Apricots

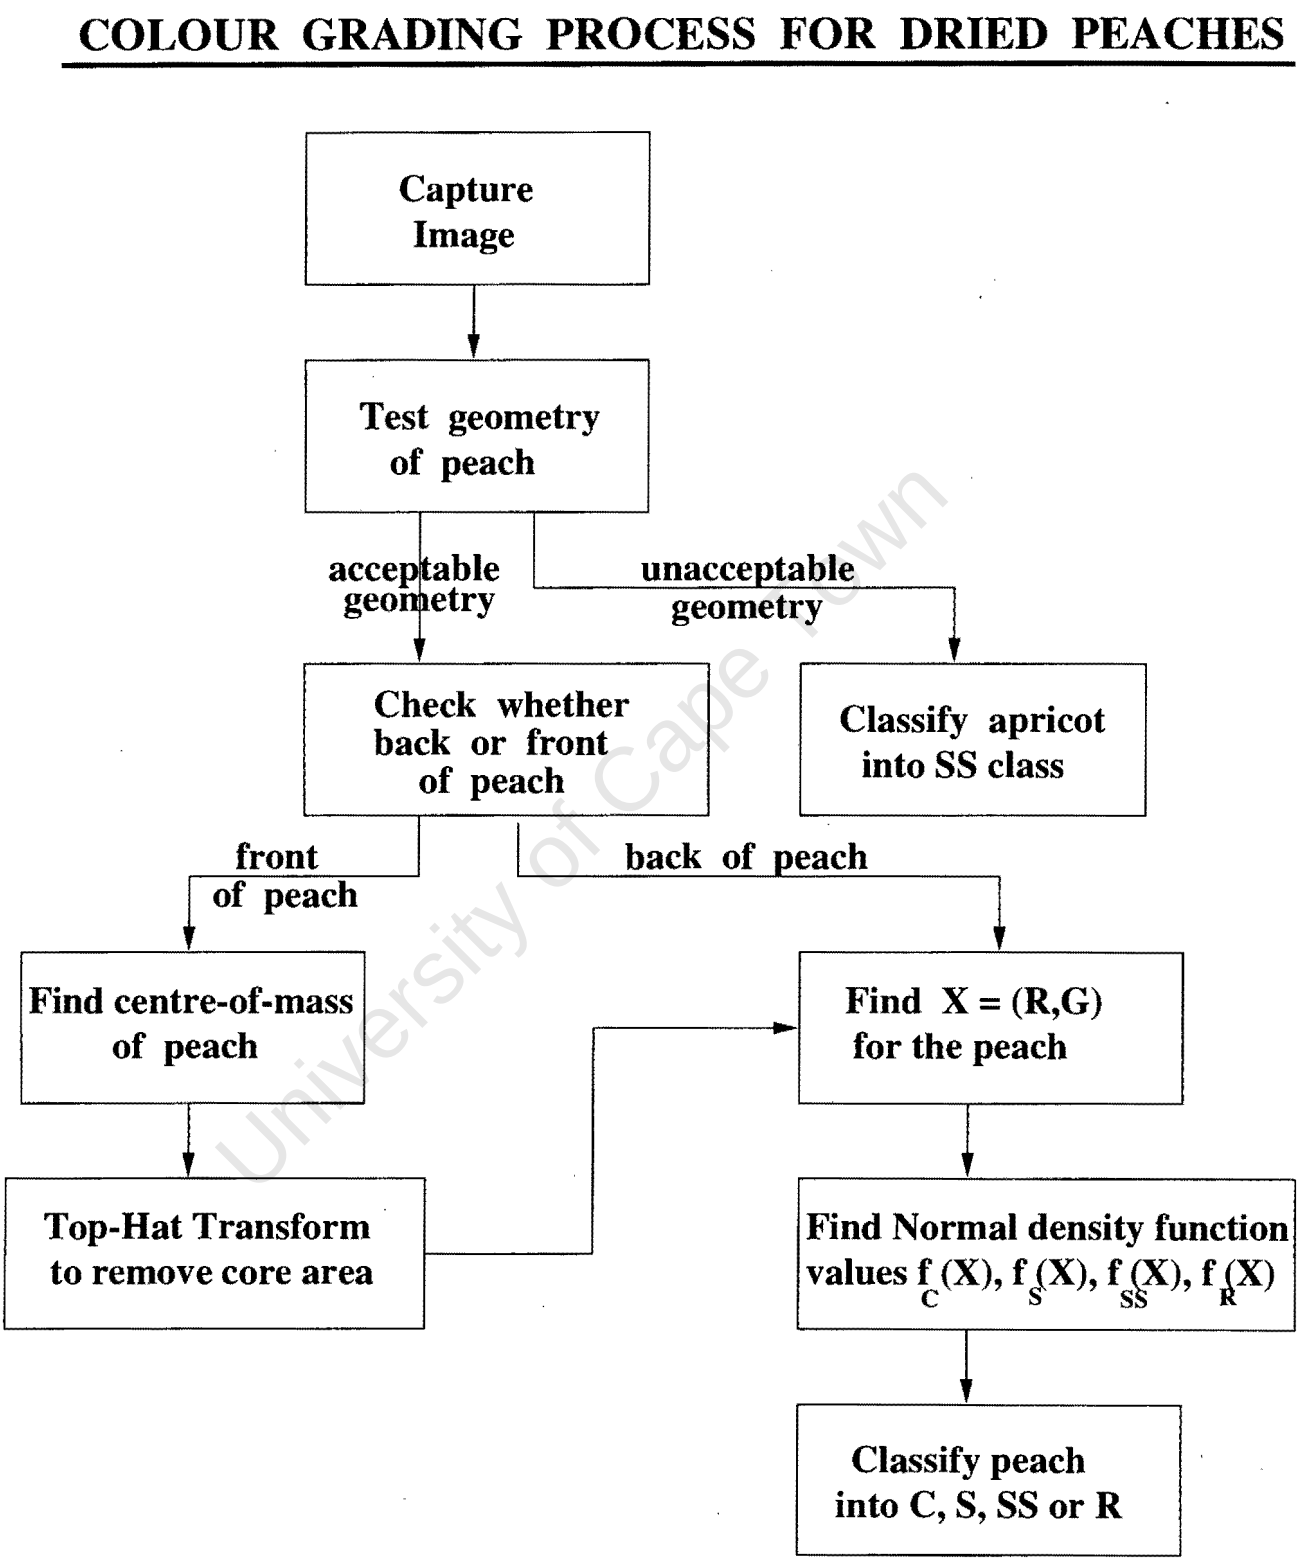

Figure 8.3: System Flowchart for classification of Dried Peaches

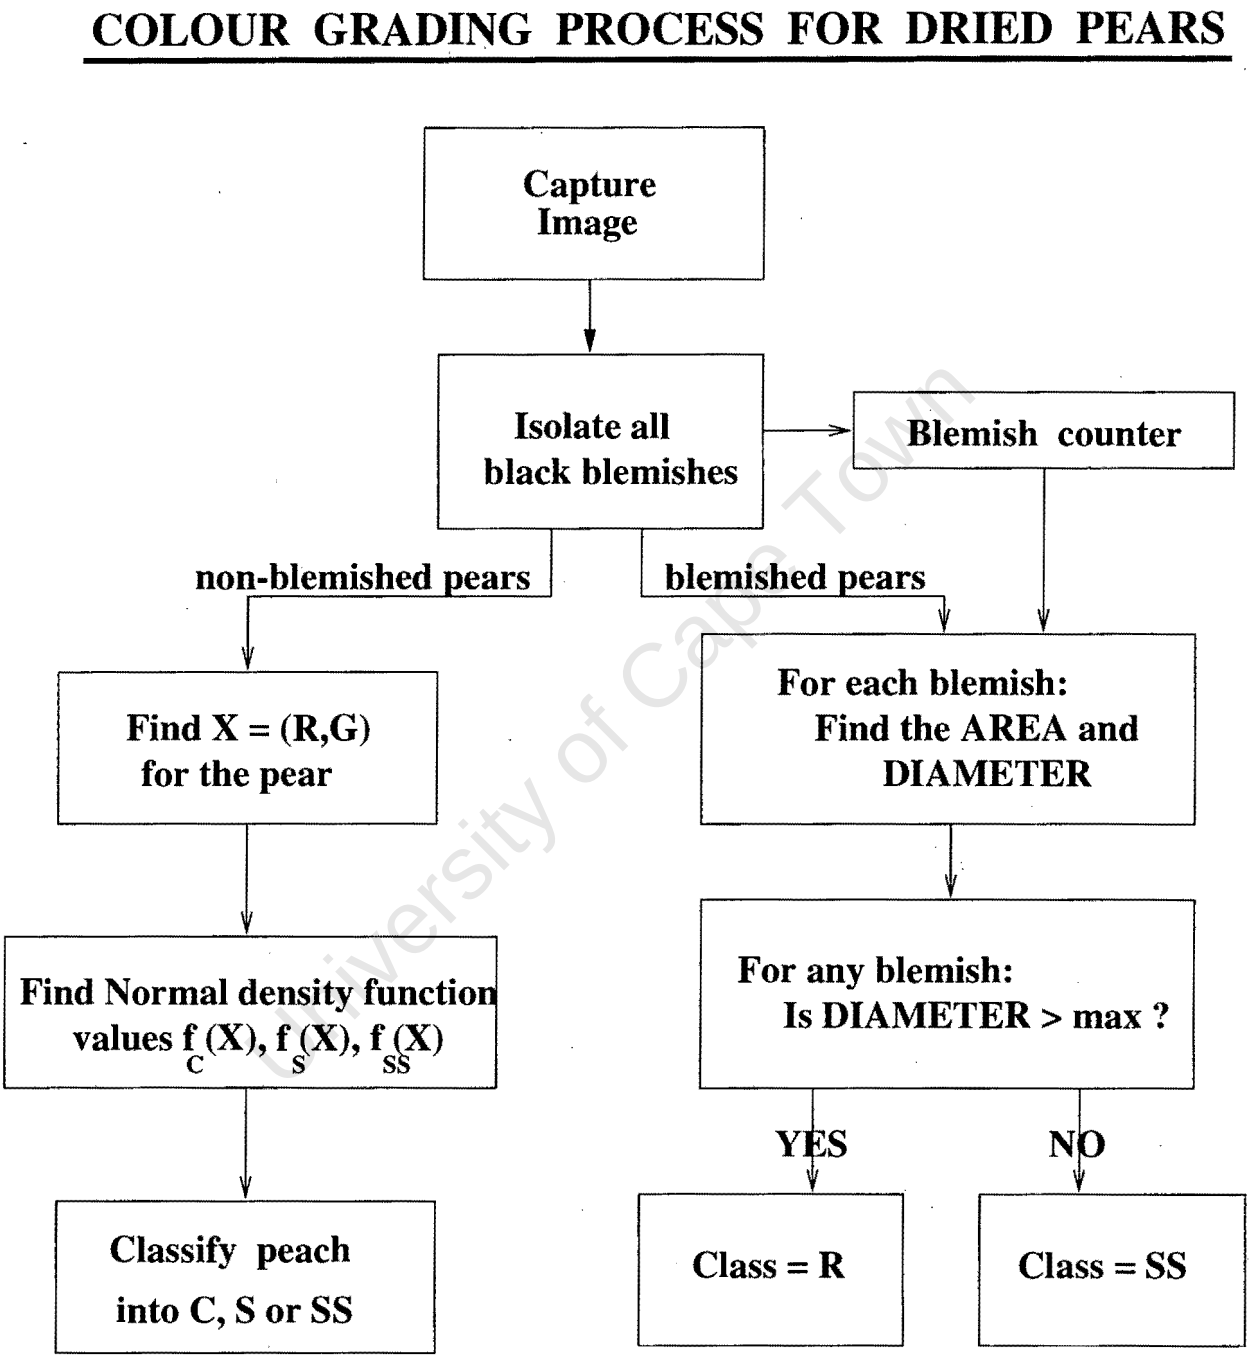

Figure 8.4: System Flowchart for classification of Dried Pears

Bibliography

[1] Gupta A. and Bajcsy R. Volumetric segmentation of range images of 3d objects using superquadric models. *CVGIP:Image Understanding*, 58:302–326, November 1993.

[2] Rosenfeld A. and Davis L.S. Image segmentation and image models. In *Proceedings of the IEEE*, volume 67, pages 149–157. The Institute of Electrical and Electronic Engineers, Inc., May 1979.

[3] Brink A.D. The selection and evaluation of grey-level thresholds applied to digital images. Master's thesis, Rhodes University, South Africa, 1987.

[4] Jain A.K. and Chen Y. Address block location using color and texture analysis. *CVGIP Image Understanding*, 60:179–190, September 1994.

[5] Miller B.K. and Delwiche M.J. A color vision system for peach grading. *Transactions of the American Society of Agricultural Engineers*, 32:1484–1490, July-August 1989.

[6] Horn B.K.P. *Robot Vision.* The MIT Electrical Engineering and Computer Science Series. McGraw-Hill, 1986.

[7] Goncalves D.P. and Barnard E. Estimating neural network pererformance using a block cross validation test. In *Fourth South African Workshop on Pattern Recognition*, pages 53–60, November 1993.

[8] Dougherty E.R. *An Introduction to Morphological Image Processing*, volume TT9 of *SPIE Tutorial Texts in Optical Engineering.* SPIE Optical Engineering Press, 1992.

[9] Kay G.R. Colour analysis and the classification of fruit. Master's thesis, University of Cape Town, Cape Town, South Africa, 1992.

[10] The Khoros Group. *Khoros Manual Volume 1: User's Manual.* Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, 1.0 edition, 1991.

BIBLIOGRAPHY

[11] The Khoros Group. *Khoros Manual Volume 2: Programmer's Manual.* Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, 1.0 edition, 1991.

[12] Spath H. *Cluster Analysis Algorithms for data reduction and classification of objects.* Computers and their Applications. Ellis Horwood Limited, 1980.

[13] Guedalia I.D., Edan Y., and Pasternak H. Multi-sensor quality classification of tomatoes. ASAE Paper No. 94-6032, ASAE, 1994.

[14] Kittler J. and Illingworth J. Minimum error thresholding. *Pattern Recognition*, 19:41–47, 1986.

[15] Liu J. and Yang Y. Multiresolution color image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16:689–700, July 1994.

[16] Tajima J. Uniform color scale applications to computer graphics. *Computer Vision, Graphics, and Image Processing*, 21:305–325, 1983.

[17] Foley J.D., Van Dam A., Feiner S.K., and Hughes J.F. *Computer Graphics - Principles and Practice.* Addison-Wesley Publishing Company, second edition, 1990.

[18] Kapur J.N., Sahoo P.K., and Wong A.K.C. A new method for gray-level picture thresholding using the entropy of the histogram. *Computer Vision, Graphics, and Image Processing*, 29:273–285, 1985.

[19] Parker J.R. Gray level thresholding in badly illuminated images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:813–819, August 1991.

[20] DaPonte J.S. and Fox M.D. Enhancement of chest radiographs with gradient operators. *IEEE Transactions on Medical Imaging*, 7:109–116, June 1988.

[21] Lim J.S. *Two-dimensional Signal and Image Processing.* Prentice Hall Signal Processing Series. P T R Prentice-Hall, Inc, 1990.

[22] Tou J.T. and Gonzalez R.C. *Pattern Recognition Principles.* Addison-Wesley Publishing Company, 1974.

[23] Fukunaga K. *Introduction to Statistical Pattern Recognition.* Electrical Science. Academic Press, 1972.

[24] Jain K., Smith S.P., and Backer E. Segmentation of muscle cell pictures: A preliminary study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2:232–241, May 1980.

[25] Rao K. and Nevatia R. Describing and segmenting scenes from imperfect and incomplete data. *CVGIP:Image Understanding*, 57:1–23, January 1993.

[26] KROMCO. Grabouw, South Africa. Personal Communication, 1996.

[27] Fu K.S. *Digital Pattern Recognition.* Springer-Verlag Berlin Heidelberg, 1976.

[28] De Floriani L. Feature extraction from boundary models of three-dimensional objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:785–798, August 1989.

[29] Celenk M. and Smith S.H. Color image segmentation by clustering and parametric-histogramming technique. In *Proceedings of 8th International Conference on Pattern Recognition*, pages 883–886, October 1986.

[30] Schwartz M. and Shaw L. *Signal Processing: Discrete Spectral Analysis, Detection, and Estimation.* McGraw-Hill International Book Company, 1975.

[31] Levine M.D. Feature extraction: A survey. *Proceedings of the IEEE*, 57:1391–1405, August 1969.

[32] Otsu N. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-9:62–66, 1979.

[33] Maragos P. and Schafer R.W. Morphological systems for multidimensional signal processing. In *Proceedings of the IEEE*, volume 78, pages 719–737. The Institute of Electrical and Electronics Engineers, April 1990.

[34] Winm P. South African Dried Fruit, worcester. Personal Communication.

[35] Robertson P.K. Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data displays. *IEEE Computer Graphics and Applications*, pages 50–63, September 1988.

[36] Mohan R. and Nevatia R. Perceptual organization for scene segmentation and description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:616–635, June 1992.

[37] Johnson R.A. and Wichern D.W. *Applied Multivariate Statistical Analysis.* Prentice-Hall Series in Statistics. Prentice-Hall, Inc., 1982.

[38] Haralick R.M. and Shapiro L.G. *Computer and Robot Vision*, volume 1. Addison-Wesley Publishing Company,Inc., 1992.

[39] Duda R.O. and Hart P.E. *Pattern Classification and Scene Analysis.* Wiley and Sons, 1973.

BIBLIOGRAPHY

[40] Hunt R.W.G. *The Reproduction of Colour.* Fountain Press, second edition, 1957.

[41] Bow S. *Pattern Recognition.* Marcel Dekker, Inc, 1984.

[42] Tominaga S. Color image segmentation using three perceptual attributes. In *Proceedings of 8th International Conference on Pattern Recognition*, pages 628–630, October 1986.

[43] Oregon Simco/Ramic Corporation, Medford. Kroma sort model.

[44] Bartleet T. Real-time automatic machine inspection of plastic bottle closures. Master's thesis, University of Cape Town, Rondebosch, Cape Town, South Africa, December 1995.

[45] Dunne T. Associate Professor in statistical sciences. University of Cape Town, South Africa, 1996.

[46] Pun T. A new method for gray-level picture thresholding using the entropy of the histogram. *Signal Processing*, 2:223–237, 1980.

[47] Uchiyama T. and Arbib M.A. Color image segmentation using competitive learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16:1197–1206, December 1994.

[48] Young T.Y and Fu K., editors. *Handbook of Pattern Recognition and Image Processing.* Academic Press, Inc., 1986.

[49] Young T.Y. and Fu K., editors. *Handbook of Pattern Recognition and Image Processing.* Academic Press, Inc., 1986.

[50] van Zyl B. Teklogic, Paarl. Personal Communication, 1991.

[51] Niblack W. *An Introduction to Digital Image Processing.* Prentice/Hall, 1986.

[52] Tsai W. Moment-preserving thresholding: A new approach. *Computer Vision, Graphics, and Image Processing*, 29:377–393, 1985.

[53] Pratt W.K. *Digital image Processing.* John Wiley and Sons, Inc., second edition, 1991.

[54] Meisel W.S. *Computer-Oriented Approaches to Pattern Recognition*, volume 83 of *Mathematics in Science and Engineering.* Academic Press, Inc, 1972.

[55] Kodratoff Y. and Hutchinson A., editors. *Machine and Human Learning.* GP Publishing, Inc., 1989.

[56] Le Grand Y. *Light, Colour and Vision.* Chapman and Hall Ltd, 1957.

[57] Ohta Y., Kanade T., and Sakai T. Color information for region segmentation. *Computer Graphics and Image Processing*, 13:222–241, 1980.

[58] Sun Y., Wu C., Lin X., and Chou N. Color image analysis for liver tissue classification. *Optical Engineering*, 32:1609–1615, July 1993.

Appendix A

Details of the Experimental Hardware

Colour images of dried fruit are required as input data to the system. The entire experimental set-up (as seen in figure A.1) used to capture and store these images was as follows :

- A JVC TK-870E RGB colour camera. This camera is equipped with a solid-state Charge-Coupled Device (CCD) image pickup element. Each image captured consists of three bands; one Red, one Green, and one Blue.
- A White background for the dried fruit samples. This choice of background colour is discussed in the Appendix.
- Sufficient lighting to provide good image colour rendering. This consisted of 4 lighting elements placed equidistantly about the camera; two 12V,50W tungsten halogen lamps and two 230V,60W incandescent light bulbs.
- The entire image capturing area was enclosed, in order to exclude ambient lighting.
- An MVP-AT colour framegrabber served to connect the RGB camera, a 16MHz 386 computer, and a Philips RGB colour monitor.

The dried and fresh fruit samples that were tested in the initial colour survey, were random samples bought at Woolworths and Pick n Pay stores. The dried fruit samples that were used to train and test the final classifier system, were all hand-picked by trained staff at South African Dried Fruit in Worcester.

Each image captured and displayed on the Philips RGB colour monitor, by the MVP-AT board used four frame buffers. Each buffer held 512 $\times$ 512 pixels. Each pixel is stored in

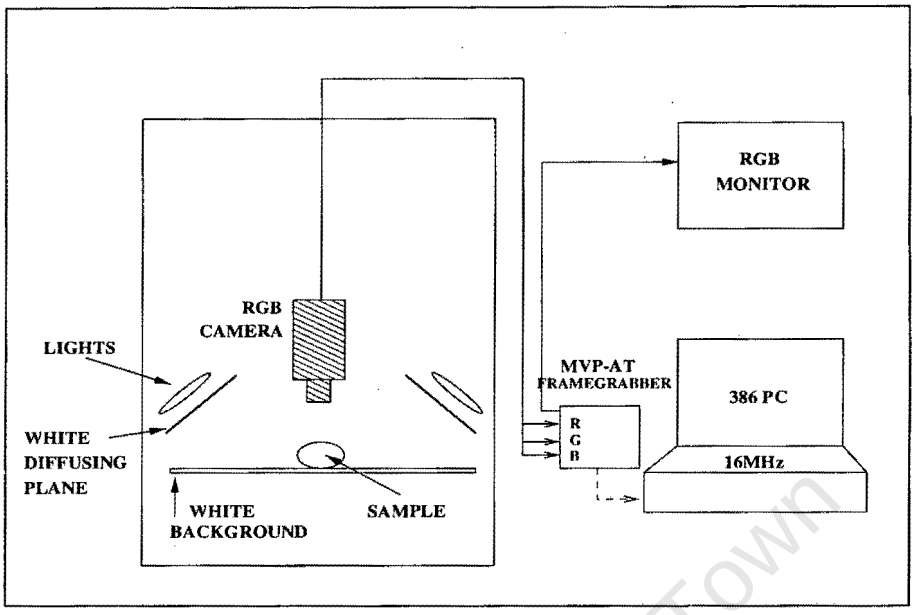

Figure A.1: Experimental Hardware configuration

an 8-bit byte implying a maximum pixel value of 255. Buffers 0,1 and 3 held the Red, Green and Blue band pixel values respectively. Combining these 3 bands produces a possible 24 bit colour value for a pixel. This implies that a range of $(256)^3$ different colour values are available to the colour monitor.

Most of the investigatory image processing procedures used, were those found in the Khoros Image Processing software package [10]. This system was installed on a configuration of Sun workstations. Algorithms that were developed for the Khoros environment were developed in the C language using the Khoros programmers manual [11].

Appendix B

Check for the front or back of a dried peach

Any image captured of a dried fruit as it moves below the camera, could be of the front or back of the fruit. In the case of dried peaches, their is a dark coloured area on the 'front' side of a peach half, where the core has been removed from the peach flesh. In order to process and grade any dried peach successfully, it is necessary to isolate this area of the peach. This process is described in Chapter 5. It is not necessary to check for core areas on the 'back' of a peach half - for this reason it would be sensible to determine which side of the peach is being processed, and further process only relevant images.

The statistical sample variance of the Red band values of 16 peach images were examined. The group consisted of :

- 4 images of the front of a Choice grade peach
- 4 images of the front of a Reject grade peach
- 4 images of the back of a Choice grade peach
- 4 images of the back of a Reject grade peach

The values for the statistical sample variance of the Red band, for each peach, were as follows :

Statistical variance values for dried peach images					
Front/Back	Grade	Variance	Front/Back	Grade	Variance
B	C	1292	F	C	8969
	C	1010		C	8512
	C	2477		C	9418
	C	529		C	8840
	R	1104		R	19969
	R	1518		R	12019
	R	1894		R	15913
	R	2828		R	15174
Back		1582	Front		12352

Variance values in the Red band are low for the back of the peach, with a maximum sample variance in the region of 3000. The variance values in the Red band for the front of the peach, are very high, indicative of the undulating, fluctuating gradient of the peach; these values tend to be larger than 8000. The variance seems to be a simple, but effective way to test whether the image is that of the front or back of a dried peach, and save unnecessary processing.

Appendix C

Finding the most appropriate colour for the image background

In order to separate the background of the dried fruit from the fruit itself, it was necessary to find a characteristic or data representation that clearly differentiated the fruit from the background. Various colour backgrounds for the 'belt'(which would carry the fruit past the camera) were examined.

Using a White or Black background for the fruit, various spatial techniques were examined on each of the following colour coordinate bands :

- Intensity band
- Red band
- Green band
- Blue band

The histograms of various fruit, as described by the above four bands, were examined. With a Black background, neither the colour of, nor any particular spatial characteristic, isolated the background entirely from the fruit. Darker or blemished areas on the fruit were often grouped with the background.

Upon examining the Red, Green and Blue colour make-up of many fresh and dried fruit, consistently low Blue values were observed, in both Choice grade and Reject samples. In every example (e.g. dried apple, banana, apricot, pear, peach) the histogram of the Blue values had clear bimodal peaks; the pixel blue values were always very low, in the range = (0..20) The colour White, on the other hand, has high Red, Green and Blue values. Examination of

a White background under various lighting conditions indicated Blue values in the region of 230. A White background was adopted as the background colour for all test images taken of dried fruit in this investigation.