THE UNIVERSITY OF HULL

SOME APPLICATIONS OF DIGITAL IMAGE PROCESSING FOR AUTOMATION IN PALYNOLOGY

being a Thesis submitted for the Degree of
Doctor of Philosophy
in the University of Hull

by

MITCHEL LANGFORD BSc

December 1988

Some Applications of Digital Image Processing for Automation in Palynolo

ABSTRACT

The work presented in this thesis attempts to solve a long standing palynological problem, the identification and counting of pollen grains, by the application of digital image processing techniques. Digital image processing has developed rapidly over the last two decades and an extensive range of techniques have been developed for the analysis of images stored in a digital format.

Exine textures, shown in great detail under the scanning electron microscope (SEM), are highly characteristic of a given pollen class. This offers great potential for automated pollen identification provided that reliable texture measures can be computed from digital SEM images.

Three methods of texture quantification, two statistically based and the other structurally based, were applied to digital samples of exine texture taken from six pollen taxa. This provided a variety of numerical texture descriptions that were used to develop statistical classification schemes.

Using a statistical classifier on the texture descriptions of unknown exine samples produced correct classification rates in the order of 70% to 98%. These success levels were not attained immediately, but were achieved by careful design modifications to a standard classification scheme to yield enhanced performance. Optimal feature selection was a major contribution to improved classification success. Combining numerical descriptions from different texture quantification schemes also produced notable improvements in classification performance.

The problem of automatically locating pollen under the SEM and selecting a suitable region for texture analysis is also considered briefly.

Acknowledgments

Although this work, including its flaws and defects, is essentially my own, I would like to thank the following people for their assistance. Full credit must go to my supervisor, Dr Gaynor Taylor, who provided support and guidance throughout the period of study, and coaxed me into completing the task. Support and enthusiastic encouragement was also given in abundance by Dr John Flenley. Paul Richards introduced me, then a complete novice, to the delights of digital image processing. His patient explanations at the start of the project are greatly appreciated. Many facilities at Leicester University were used to produce this entirely computer generated report. My present employers have been kind enough to allow me to work on its presentation during slack moments (thanks especially to David Unwin). Finally, and by no means least, I have to say a "chears-big-ears" to Andy, Dan, and Nigel, who ensured that the three years of study were also a great deal of fun.

CONTENTS

CHAPTER ONE

INTRODUCTION

1.1 INTRODUCTION

Palynology and digital image analysis are two fields of study which have until now had little, if indeed any, connection. The work presented here draws these two diverse sciences closer together as it attempts to solve a long standing palynological problem, the identification and counting of pollen grains, by the application of modern digital image processing techniques.

Before specifying the nature of the identification problem in detail, a brief résumé of palynology that focuses particularly on pollen analysis is presented, for the sake of those unfamiliar with the principles involved in this subject. The need for a solution and its value to the palynologist can be more fully appreciated once a basic understanding of these principles has been established.

1.2 PALYNOLOGY

The term micropalaeontology is used within the science of geology to refer to the study of microfossils. We can consider palynology to be a branch of micropalaeontology that is concerned specifically with the study of pollen grains and spores. Pollen grains are the male reproductive organs of the flowering plants, known as the angiosperms. Similarly, spores are the reproductive bodies of fungi and ferns, known as the gymnosperms.

In fact palynology is a rather broader multi-disciplinary field than the simple definition given above would suggest. It has, for example, found applications in the exploration of petroleum and coal deposits, and in

studies of Quaternary history, archaeology, criminology, and medicine.

Palynologists are concerned with all aspects of pollen and spores. They study their formation and morphology. The processes of dispersion from the parent plants and the potential for preservation in different natural environments are investigated. However, perhaps the most important and useful area of study is that of pollen analysis. Mannion (1980) provides a simple introductory guide to the theory and applications of pollen analysis. In order to understand the fundamental principles involved it is first necessary to identify some important characteristics of pollen itself.

1.3 POLLEN AND POLLEN ANALYSIS

Plants are capable of producing pollen in enormous quantities, especially those that rely on wind pollination. For example, it has been estimated that over 40000 grains may be released from a single anther of the Pine tree (Moore and Webb, 1983), and there may be hundreds of such anthers on any one tree.

Pollen grains are very small, typically they have a mean diameter of between $1.5x10^{-7}$m and $5x10^{-7}$m, and as a consequence they are extremely light which facilitates their rapid and widespread dispersal by wind. This in turn ensures that the pollen released from plant species within a local area become thoroughly and homogeneously mixed before eventually settling out as a 'pollen rain'.

With vast production by local plant species and thorough mixing and dispersal, the pollen settling out from the atmosphere should be an

accurate reflection of the plant population of the region. Pollen analysis relies on the assumption that if, at any given location, no pollen of a species settles, there is minimal chance of that species being present within the local area (Faegri and Iversen, 1975).

Another important feature of pollen is that the grains have a much greater potential for preservation than almost all other plant materials. This fact can be attributed to the construction of the cell wall which is composed of two layers. The inner layer, the intrine, is a normal cell wall built from cellulose. However, the outer wall, the exine, is composed of a material called sporopollenin. This is a very complex natural substance that has an extraordinary high resistance to both biological and chemical attack. Biological enzymes and many chemicals, particularly acids, have little effect on its decomposition. Therefore, any pollen that settle onto an actively accreting surface, such as a peat bog or lake bed, have a good chance of preservation as they become buried. This is particularly so if anaerobic (oxygen deficient) or acidic conditions exit. Eventually, over geological time spans, they may withstand the heat and pressures associated with diagenesis (rock formation) to become fully fossilized remains.

A final feature, and one that is of fundamental importance, is that despite their microscopic size pollen grains are highly recognizable objects. Frequently it is possible to discriminate between them down to the plant species level. All the information needed for discrimination is contained in the morphology of the exine. This last feature allows pollen that has been extracted from a sediment to be identified and counted. These proportions may then be related back to a vegetation assemblage.

This then leads us to the fundamental principle of pollen analysis: **The assemblage of pollen contained in a sediment has a relationship to the assemblage of plants in the local area at the time that the sediment was deposited.**

Pollen analysis consists of the quantitative evaluation of pollen assemblages extracted from successive levels of a sedimentary core (Brasier, 1980), allowing the reconstruction of the vegetation history of a region. The cores from which pollen are extracted are typically of peat bog or lake bed sediments, but they may also be older geological deposits such as sandstones and clays.

Of course, a preserved pollen assemblage is a biased estimate of the ancient plant population. It is necessary to take into account the production rate, the dispersal characteristics, and the preservation potential of different species. However, with these borne in mind it is possible to reveal remarkable details of vegetation history within a region.

1.4 APPLICATIONS OF POLLEN ANALYSIS

The true power of pollen analysis depends not only on the ability to decode vegetation history but also on the insight that this gives us of the natural environment. The vegetation found within a region is closely linked to the regional climate and other environmental factors. By analogy with present day vegetation patterns a pollen assemblage can be used to shed light on the climatic and environmental history of an area. The value of pollen analysis is easily demonstrated with reference to studies of the geological Quaternary Period (0-2 million years ago).

Pollen analysis of Quaternary sediments from northern Europe has revealed cyclical variations in the vegetation assemblage. Each cycle can be attributed to the climatic changes occurring during an interglacial period (an interval between the 'ice ages'). It has been possible to define a number of pollen-analytical zones, each identified by a characteristic vegetation sequence. This in turn has enable pollen analysis to become an important method of relative dating, and has made wide correlations between sediments possible over the greater part of northern Europe (West, 1977).

Archaeologist have been frequent users of the pollen zone dating system. Pollen analysis can also provide them with information regarding the impact of early man on the environment. The first acts of deforestation, for example, are often clearly identifiable in the pollen record.

The work of Van der Hammen et al. (1973) is a spectacular example of the application of pollen analysis. Their investigation of Pliocene and early Quaternary sediments in the Columbian Eastern Cordillera produced an astonishing palynological record of the upheaval of the Andes mountain range. The pollen assemblages showed gradual changes in the vegetation from tropical lowland flora, through intermediaries, to high mountain flora. From the data obtained it was possible to estimate the average rate and the period of uplift of the mountain range.

Fossilized pollen are found in rocks dating as far back as the Devonian Period (approximately 390 million years ago). They are frequently contained in rock strata that are closely related to deposits of crude oil and coal. In these rocks other fossils are often scarce and pollen may then become essential stratigraphic markers that aid the correlation of

the order of three or four hundred are ever likely to be found in any one sample. Of these the vast majority will be very rare and appear only occasionally. In a typical northern hemisphere sample only twenty or thirty taxa will occur in any significant number. Furthermore, it is possible to derive useful palynological information from the relative proportions of only a few, say six to ten, common taxa. Therefore the identification task is considerably simpler than it at first appears. A system that could automatically identify and count only the most common taxa in a sample, and express these as a proportion of the total pollen content, would be a very powerful palynological tool.

All pollen identification and counting is currently performed manually. Human powers of visual recognition and interpretation are heavily relied upon. By employing techniques from digital image analysis it may be possible to provide an alternative to this situation.

Due to the size of pollen a high powered microscope is an essential requirement for viewing. Standard optical microscopy is typically employed but occasionally, and ever more frequently, use is made of the scanning electron microscope (often abbreviated to SEM).

The first stage in the identification process is to scan the sample in order to locate pollen, which are often situated among other artifacts left over from the preparation process. This search is carried out using a special microscope stage that allows a grid-like movement of the sample in order to prevent the repeated identification and counting of any individual specimen.

Once a pollen grain is located it is studied carefully by the analyst and a decision made on its identity. The morphological features of the

exine provide all the information necessary for its identification. The specific features employed in the identification process depend on the specimen itself and to a large extent on the type of microscope in use. A running total is kept of each pollen taxa found in the sample and these raw data are passed on to the next major stage of the process.

It is not possible to specify the exact time taken to analyse a sample since this obviously depends on many factors. A crude idea can be given, but it must be considered as such. To produce results that have some statistical significance the identity of at least 200 grains are generally considered necessary. Normally somewhere between 200 and 400 grains would be analysed but in some situations it may be necessary to identify many more, 1500 grains is not uncommon. The quality of the preparation, the experience of the palynologist, the variety of pollen types present, and the familiarity of the analyst with these types, are all factors that will affect the time taken for the analysis. However, suppose a count of 400 grains was required for a 'typical' sample, one in which the majority of the pollen present were familiar to the analyst. A palynologist of average ability would not expect to complete this task in under two days.

The majority of palynologists, most of whom are graduates, regard the identification and counting of pollen as a burdensome chore. It is the interpretation of the data that provides interest in the work. Clearly then an automated identification system would be very welcome.

1.6 DIGITAL IMAGE PROCESSING AS A POSSIBLE SOLUTION

Digital image processing has developed rapidly since the early 1970s and continues to do so today. Over the last two decades an extensive range of techniques have been developed for the manipulation and analysis of images stored in a digital format. Thresholding, edge detection, and texture analysis are three examples that will be seen later in this report. By suitably combining these basic techniques it might be possible to construct a system that will allow the automatic identification and counting of pollen.

Many identification systems employing digital image analysis already exist. These cover a wide range of applications but the majority are designed to deal with objects of well defined form. The identification of components on a production line is a typical example. There are, however, some systems working with biological materials. For example, Rutovitz et al. (1978) describe the application of image processing for the automated analysis of chromosomes. A commercial system that classifies and counts white blood cells is currently in use in several British hospitals. There have as yet been no attempts to analyse pollen using the digital image processing approach and this material presents a number of unique problems to overcome.

The term pattern recognition is often used to describe the application of computers to recognition problems. This covers a very broad range of applications such as speech interpretation, the reading of hand written characters, and the location of objects within a scene. In all pattern recognition problems three fundamental stages can be identified.

The first stage, and one that is not necessary when the task is

performed manually, is the digitization of the data. This means that some form of relevant data must be captured and converted into a numerical format for storage and analysis by the computer. For a pollen recognition system this does not present a major difficulty as we simply require digitally encoded images of pollen taken from under the microscope. This can be readily achieved with suitable hardware, either attached to the microscope itself, or via a video scanner and photographic prints.

Once a digital description of the object has been stored the stage of feature extraction is needed. This is equivalent to an analyst recognizing the shape of a grain when identifying pollen manually, or to a computer counting the number of sides on a object when searching for a specific part on a production line. It is at this stage that digital image processing and analysis will play a major role in automated pollen identification.

The final stage is to act upon the data provided by the feature extraction procedures. For pollen identification we need to use the data to classify the object. Statistical classification schemes will be required to perform this task.

1.7 AN OUTLINE OF THE THESIS STRUCTURE

In the following chapter a review of the literature citing previous relevant work is presented. It is divided into two sections, the first on palynological aspects and the second on digital image processing techniques. The palynology section outlines recent advances in preparation techniques that simplify the identification problem. It also

discusses previous work aimed at providing automation in pollen analysis. In the second section a review of selected items from the digital image processing literature is presented. This selection concentrates on the techniques that are most likely to prove useful in solving the pollen identification problem.

In Chapter Three details are given of the equipment and practices used in order to obtain and process the digital image data. This includes information on the source of pollen and the preparation techniques used, as well as details about the hardware and software used to capture and analyse the images.

The remainder of the report broadly follows the sequence of stages identified above for all pattern recognition problems. Chapter Four deals with the image processing procedures that are required to extract or measure useful features from the scenes of pollen available.

Chapter Five outlines the various statistical classification schemes available for identifying the objects from the extracted feature information. It also describes some variable selection procedures that may be used to enhance the performance of a classifier.

In Chapter Six a summary of experimental results is presented. These were obtained by using the procedures covered in the earlier chapters. Chapter Seven looks very briefly at some possible approaches to the problem of locating pollen within a given scene. Finally, Chapter Eight draws conclusions from the work currently undertaken and identifies areas that are likely to prove profitable for future development in this field of study.

CHAPTER TWO

LITERATURE REVIEW

PALYNOLOGY

2.1 POLLEN PREPARATION TECHNIQUES

The extraction of fossilized pollen from a sediment and its preparation prior to viewing under a microscope are both important aspects of palynology. The selection of a suitable preparation process is a vital consideration since pollen can react differently to alternative methods, and the visual effects often vary between taxa (Cusma Velari, 1984). The choice of preparation technique becomes more critical if automated identification is required since features such as the shape, size, and the surface detail of the grain must be kept as constant as possible for effective machine interpretation.

2.1.1 Extraction of Pollen from Sediments

An extensive range of methods is available to deal with the removal of sediments in which pollen is contained. Most are based on chemical processes and they are adequately described in the standard palynological texts (e.g. Faegri and Iversen (1975), Moore and Webb (1983)).

The process of acetolysis deserves a special mention since it is used as the first step in most preparation techniques. This consists of immersing the sample is a series of increasingly powerful acidic solutions in order, primarily, to remove the inner cellulose layers of the pollen. The sample is then passed back through a second sequence of progressively weaker solutions until it is neutralised.

After treatment with a series of suitable chemical processes a typical core derived sample still contains a considerable quantity of material other than pollen. This often consists predominantly of amorphous organic materials and silica spicules. When manually identifying the grains these impurities are usually left in the sample unless the problem is very severe. However, for an automated system it is essential that as clean a preparation as possible is obtained in order to minimise the complexities of locating pollen. The unwanted artifacts are most readily removed by recently developed physical extraction and concentration techniques.

By the use of density gradient and centrifugation equipment an extremely clean, pollen rich sample may be obtained (Forster, 1986). The separation of some groups of pollen taxa is also possible. Although this is a difficult process to use at present, it does have potential for future automation.

Ultrasonic sieving or filtration (Caratini, 1981) can also produce a considerable improvement in the quality of prepared samples. Ultrasonic shaking of the sample on extremely fine mesh sieves removes material outside the size range of pollen grains. This technique has been found to be particularly effective on silt and clay rich deposits which are difficult to process chemically without causing damage to the pollen (Tomlinson, 1984). An example of the combined use of chemical and physical extraction techniques in order to produce usable samples is presented by Heusser and Stock (1984).

2.1.2 SEM Preparation Techniques

The advantages offered by the scanning electron microscope (SEM) over conventional optical microscopy include much greater magnification and definition and a remarkable depth of field. However, the environment within the SEM is much harsher on the samples. Specimens must be completely desiccated and are subjected to a high vacuum during viewing. Pollen is relatively rigid and indestructible compared to most other biological materials and yet serious problems can still occur with many taxa (Adams and Morton, 1972). The most important of these are grain distortion, excessive contrast, and charging (the build up of electrons on the surface which deflects the electron beam). Since the scanning electron microscope became available to palynologists a series of new preparation techniques have been developed, and many of the traditional methods have been modified.

Initially, the preparation techniques of optical microscopy were used for SEM samples. The only alteration was the use of a conductive carbon or metal film to prevent charging of the specimen. This film was applied by an evaporation process. Adams and Morton (1972) proposed a greatly improved procedure which displayed two essential features. Firstly, acetolysis was used to reduce the contrast and glare in the image. More importantly a critical point drying method was employed to prevent the distortion and collapse of the grains. Collapse is normally caused by surface tension forces generated during drying, but these are eliminated by the critical point drying process. They also recommended the use of gold or aluminium as the coating material as this helped to further reduce image contrast. However, an aluminium coating has the disadvantage that it deteriorates rapidly, so a sample must be analysed shortly after its preparation.

The advantages of critical point drying over conventional methods have been extolled in many later papers. For instance, Nilsson et al. (1974) noted consistently better results when comparing this method to the vacuum and excess-pressure treatments used previously. Cusma Velari (1984) found critical point drying and dehydration through acetone solutions to give the most satisfactory results in a comparison of five preparation methods.

The virtues of acetolysis in SEM preparation have been the cause of some debate. Rowley (1973) suggests acetolysis limits fine detail and causes distortion and shrinkage of the grains. Lynch and Webster (1975) point out that acetolysis is often a disadvantage for fresh pollen since removal of the inner cellulose layer encourages the grains to collapse.

A further significant advance in SEM preparation was presented by Damblon (1975). He introduced the sputtering technique as a means of applying the conductive coating to a sample. A metal target, typically gold, is bombarded with ions under a low vacuum. This causes metal atoms to be ejected onto the specimen producing a thinner and more even coating than the evaporation method. This technique improves the level of surface detail and the consistency of image quality. It has now been adopted almost universally.

2.2 AUTOMATION IN PALYNOLOGY

Previous work at providing automation in palynology has been notably lacking. The analysis of raw pollen taxa frequency data by computer to provide statistical interpretations has been the most obvious development.

A number of packages are now available to achieve this (e.g. Hill, 1979a and 1979b) and many include facilities for the production of pollen diagrams or other graphical presentations of output (e.g. Squires, 1970).

There have also been some attempts at automating the preparation processes. These range from simple mechanical manipulators to improve the mounting of pollen onto an SEM stub (Leffingwell and Hodkin, 1971), to the development of sophisticated extraction and purification techniques (Forster, 1986).

Walker et al. (1968) describe a computerized pollen database system. This was designed primarily to assist palynologists in the identification of unfamiliar tropical and southern hemisphere pollen of Quaternary age. The morphological characteristics of an unknown grain are encoded and compared to the coded descriptions of modern reference pollen. The coding is divided into a number of sections, each section dealing with one of the major morphological features used in normal manual identification. The pollen taxa with the highest number of section matches are listed as the output. These are treated as suggestions which may be investigated further by traditional means.

An improved version of the original system is described by Guppy et al. (1973). The number of descriptive sections was increased and many features were defined more precisely. In addition, they incorporated greater sophistication into the presentation of printed output. The extention of the system to include fully fossilized pollen was discussed, and they reported that a number of early Tertiary grains had been successfully encoded.

A new system for the description and coding of pollen and spores was proposed by Germeraad and Muller (1970, and 1971). This attempted to remove the subjectivity of the scheme presently adopted, and is inherently more suitable for computer manipulation in a database environment. The construction of a data bank employing the new numerical coding system was proposed. This could be used, like the Walker system, for the identification of unknown pollen. However, so far the new classification scheme has not found great success. The less precise, but much simpler, traditional system is still preferred by the majority of palynologists.

An early attempt at the identification of palynological objects using image based data is presented by Mirkin and Bagdasaryan (1972). They record that "...the memory storage of modern computers is very limited.", a sharp reminder of the enormous progress made in recent years in digital computing power. Due to this limitation they concentrated on an optical based system to produce Fourier power spectra from photographic images of pollen and spores taken under a standard optical microscope. They then employed a correlation measure in order to match these images with Fourier spectra from type specimens. The results were not conclusive, but they noted that the "...microelements of the image play an appreciable role in the recognition of objects...". Thus suggesting that edges and textures are more valuable than broad tone levels and shape in the identification of the objects.

Dickson et al. (1977) describe a similar approach for the identification and counting of diatom communities in an attempt to provide automated ecological monitoring. They, like Mirkin and Bagdasaryan, produced Fourier power spectra using laser light projected through negatives of

diatoms taken from under an optical microscope. These were filtered using Fourier spectra obtained from type specimens and the image reformed to give bright spots of light on an output plane whenever close matches occurred. The spots were counted and the process repeated with alternative filters to identify other species present.

Dickson et al. noted several difficulties with their system. The size of diatoms, like pollen, vary and a variation of only 4% could reduce the signal-to-noise ratio of spots on the output plane by 50%. Sample orientation was also a problem, requiring the specimen to be rotated and several attempts made to count correlation spots. The optical microcopy gave very limited field of view so that invariably only a few diatoms were in focus at any one time. This was a big problem when using negatives, but even a direct output from the microscope would require some sort of automatic focussing system which is not an easy task. Finally, the alignment of the whole system was critical. It had to be kept within $2m^{-6}$ using sensitive step motors under computer control.

Case et al. (1978) simplified the procedure for diatom identification and counting. A laser beam was used as the source of illumination in the microscope to produce coherent images directly, without the need for the photographic stage. Several optical filtering operation were required, however, in order to match the image quality obtained from the previous method using film negatives.

DIGITAL IMAGE PROCESSING

2.3 INTRODUCTION

It is possible to identify two basic aims of digital image processing (Gonzalez and Wintz, 1977). The first of these is to improve an image in some way such that it is clearer, easier, or contains more information, for human perception. The second is to process image data by machine, almost invariably a computer, for automatic interpretation.

Sometimes a process or technique can accomplish both aims simultaneously, but more often it will assist in only one. For example, a smoothing operation to reduce noise can make a scene more pleasing for a human observer and at the same time assist an automated system when segmenting the image into meaningful regions. On the other hand colour coding, the substitution of colours for grey tones, is an example of the second case. A computer will deal with numerical image data regardless of whether it represents, or is represented by, grey tone or colours. However, the human visual system has much greater powers of discrimination and information extraction when dealing with colours than with monochrome shades. Colour coding can bring out information or detail which, although present in the original image, is otherwise lost to a human observer.

2.3.1 Brief History

An interesting account of the early development of digital image processing is given by McFarlane (1972). One of the first applications, around the turn of the century, was the transmission of digitized newspaper pictures between London and New York via a submarine cable. In 1929 the tonal resolution of this system was improved from 5 to 15 grey levels. Similar small advances in digital imagery were made up until the early 1960s when the space program and the increasing availability of powerful digital computers provided a spark for massive development.

Gonzalez and Wintz (1977) document the work of the Jet Propulsion Laboratories in California where many basic digital image analysis techniques were developed. These were used to correct the various distortions found on digital images received from satellites such as the Ranger and the Mariner series. By the early 1970s attention was turning to other applications and the Jet Propulsion Laboratories themselves reported work on such diverse topics as improving forensic fingerprint images, and the removal of atmospheric distortions from astronomical images (O'Handley and Green, 1972). The expansion of applications still continues to accelerate today. Digital image processing now finds its way into an enormously broad range of subjects and is particularly well founded in the fields of biomedicine, computer vision, astrophysics, and Earth resource sciences.

2.4 IMAGE PROCESSING TECHNIQUES

Understandably, with such an enormously diverse range of applications, the number of techniques developed for processing digital images is extremely large. Some of the methods are now discussed below. By no means should this be considered an exhaustive discussion of image processing techniques, since it aims to concentrate on those that have particular relevance to solving the problems presented by automated pollen identification.

2.4.1 Point Operations

Point operations are a simple but important class of image processing technique. By definition a point operation produces an image in which the value of each individual output pixel depends only on the intensity level of the corresponding input pixel. The relationship between the input and output levels is defined by a transformation or 'mapping' function (Castleman, 1979).

Point operations are particularly important for modifying the way in which the image data fills the grey level range available. Point operations modify the range of grey scales in the image, and consequently modify the grey level histogram (the frequency distribution of grey levels) in a predictable manner.

Linear point operations stretch, compress, or shift the distribution within the histogram. Linear contrast stretching is a typical example in which the grey level range of the histogram is stretched so that it fills the full range available, assuming this was not achieved in the original

image. The details of this and other grey scale transformation functions are adequately covered in standard digital image processing texts (e.g. Hall (1979), Gonzalez and Wintz (1977), Rosenfeld and Kak (1976)), as well as many of the remote sensing texts (e.g. Curran (1985), Lillesand and Kiefer (1987)).

An important point operation is that known as histogram flattening, histogram equalization, or equal probability grey level quantization. This process aims to produce an image that has an equal number of pixels in each grey level bin (hence giving a flat histogram). This is particularly valuable since it allows the normalization or standardization of the first-order statistics of the image (such as the mean and variance of the grey levels). It is used as an early preparation stage of many texture analysis procedures. Castleman (1979), Niblack (1986), and Haralick et al. (1973) give details on how the transformation is implemented. More complex algorithms based on this theme are also available, such as that proposed by Alparslan and Ince (1981) who employ a locally adapted histogram stretching scheme. These are particularly suitable for retaining maximum information when using low tonal resolution output devices, such as the majority of printers in current use.

Peleg (1978) describes an iterative histogram modification scheme based on a similar idea by Rosenfeld and Davis (1978). Grey level frequencies are transferred to large bins from nearby smaller bins resulting in a modified histogram in which only a few spikes remain. Hence, this process is often called the 'super-spike' algorithm. Frequently the image corresponding to the modified histogram is barely distinguishable from the original. Image segmentation and data compression techniques such as run-length encoding or quadtrees (Walker

and Grant, 1986), are more easily facilitated with the image in this modified form.

2.4.2 Thresholding

In many applications images contain objects which are lighter or darker than the background on which they occur. Thresholding is a technique that is commonly employed to segment such images into their object and background points since it is a relatively simple and fast procedure. Although often classified as a segmentation technique, thresholding may also be considered as a type of point operation, since the output pixel values depend only on their corresponding input values.

Thresholding is performed with reference to the grey level histogram of the image. We assume that the histogram displays a bimodal distribution with the two peaks corresponding to the grey level populations of the objects and background points. Between these peaks should lie a valley of less frequent intermediate grey levels which normally correspond to points lying on the boundaries between the objects and background (Ahuja and Rosenfeld (1978), Kirby and Rosenfeld (1979)). A threshold value is selected from within this valley and the image is segmented into a binary format by mapping all points below the threshold to black and all those above to white, or vice versa.

However, the selection of a suitable threshold is frequently a non-trivial task. The segmentation of textured regions presents one problem, which will be discussed later, but even with uniform regions a valley in the histogram is often difficult to detect (Weszka et al. (1974), Otsu

(1979)). This situation arises especially when the proportions of object and background are very dissimilar (Weszka et al. (1974), Kittler and Illingworth (1985)). The histogram may then become predominantly unimodal in character with one side of the peak displaying a shoulder or a change of slope, or it may simply show a skewed distribution (Weszka and Rosenfeld, 1979).

2.4.3 Histogram Modification To Aid Thresholding

There have been a number of attempts to modify or enhance the grey level histogram in order to make the selection of a suitable threshold easier. These methods take into account not only the grey level at each point but also local property values, in particular the rate of grey level change, in order to produce an enhanced histogram. With enhancement the valley is deepened or the histogram is converted to display a sharp peak near the optimal threshold level. Reviews of the various techniques developed are given by Weszka (1978), and Weszka and Rosenfeld (1979).

For example, each point can be given a weighting, depending on its edge strength value, before contributing to the histogram. Similarly the Laplacian operator can be used as an indicator of grey level gradient (Weszka et al., 1974). Those points with a high output, normally associated with an object boundary location, may be discarded to deepen the valley between the background and object populations. Alternatively, if those with a low output are discarded, a bimodal histogram is produced whose sharp peaks are on either side of, and close to, the ideal threshold value.

The use of second-order statistics from co-occurrence matrices are an alternative method of distinguishing between border points and those lying in the interior of object or background regions. This appears to give better results than using edge values, particularly when selecting only the border points to produce a strongly peaked histogram (Ahuja and Rosenfeld, 1978).

Yet another approach is based on the use of scatter plots in {grey level/local average grey level} space (Kirby and Rosenfeld, 1979). As with the co-occurrence matrices, entries near the main diagonal represent interior points, while the distal entries should represent border points. This method, that closely resembles the Laplacian based approach, is no more expensive to compute than the second-order statistics and may be particularly valuable in cases where the locally averaged image is already being computed for other purposes (e.g. noise reduction).

2.4.4 Automatic Threshold Selection.

None of the transformation processes discussed above address the problem of automatically selecting an optimal threshold. They simply enhance the histogram to make manual selection easier. Detecting the two peaks and locating a suitable threshold level may be difficult even after enhancement, especially when the histogram is noisy. In many applications of digital image analysis unsupervised operation is required, so automatic threshold selection is essential.

One approach to automatic selection is to use the 'busyness' measure of a co-occurrence matrix constructed from the original image (Weszka and

Rosenfeld, 1978). A threshold level is selected which minimizes this measure. An alternative is the iterative selection method outlined by Ridler and Calvard (1978) which yields successively refined threshold positions. These converge to an optimum in about four iterations. This process effectively minimizes the within-class grey level variance of the object and background regions (Kittler and Illingworth, 1985). A mathematically efficient implementation of the method, which processes the histogram rather than the image itself, is described by Trussel (1979).

Otsu (1979) approached the task of automated thresholding from the viewpoint of discriminant analysis. He proposed a discriminant criterion measure which is a determinant of the between-class separability. The grey level giving the highest output is taken as the optimal threshold point. Essentially this is the reverse operation to that proposed by Ridler and Calvard (1978), but Otsu suggests that his method is simpler requiring the computation of first-order rather than second-order moments. Importantly it is a non-parametric technique so it has general applicability, and is ideal for use in an unsupervised environment.

More recently, Kittler and Illingworth (1985) have shown Otsu's method to break down for certain ratios of object and background areas. The assumption of unimodality for the discriminant criterion does not always hold, and it can become bimodal or even multimodal. A local peak determination and simple 'valley check' are required to ensure the correct threshold level is selected. This process can also usefully indicate when a single threshold and the resulting binary image are inappropriate classifications of the image.

2.4.5 Thresholding Of Textured Regions

Simple thresholding does not work effectively for textured regions. Although regions may have different average grey level, a fixed threshold level between these averages will still lead to misclassification of points due to the range of grey tones within the textured regions. Davis et al. (1975) showed that a fixed smoothing method can allow such regions to be separated. The size of the averaging neighbourhood needed for effective separation is difficult to predict however, and an experimental method of automatic variable averaging did not prove to be successful.

The smoothing approach can be generalized to cases where the average of any local property, other than simple grey level intensity, varies between the regions. The use of a local digital gradient and the Laplacian operator were demonstrated, but other more specialized operators could be devised for separating specific textural regions.

One serious problem is that this method will fail under certain conditions. Specifically, if there are three or more regions to be separated, each with distinct local property value ranges, and regions with high and low ranges are adjacent, the methods proposed above will not perform satisfactorily.

Tomita and Tsuji (1977) provide a solution to the multiple region problem with an improved smoothing method. At each point a gradient operator is used to calculate non-homogeneity indices from each of four surrounding local regions. The averaging procedure is applied to the region with the lowest non-homogeneity index. This should prevent averaging from being performed over areas that contain a boundary

between regions. The smoothing process is iterated until distinguishable regions are obtained.

Nagao and Matsuyama (1979) believe that the rectangular smoothing window employed by Tomita and Tsuji still leads to blurring and does not yield good results when applied to complexly shaped regions. They propose smoothing within an elongated bar that is rotated around each point, allowing nine positions in a standard raster-type digital image. Again smoothing is performed within the region with the lowest non-homogeneity index. This procedure ensures that smoothing occurs without the incorporation of edges, even in complex shapes. Thus the textured regions can be smoothed without losing definition of shapes or edges. Iteration of the process gives an effect equivalent to averaging over a larger neighbourhood. It was also demonstrated that the technique also has the useful ability to sharpen initially blurred object edges.

2.4.6 Local Operators

The use of local operators or spatial filters is extremely common in image restoration and enhancement applications. Castleman (1979) defines a local operator as a filter whose output at a pixel is a function of the input values within the neighbourhood of that pixel. Schowengerdt (1983) notes that spatial filtering is therefore in effect a 'context-dependent' operation. The neighbourhood is often referred to as a window. The window may be any shape, although square is most common, and almost all are symmetrical around a centre pixel. The window is scanned exhaustively across the image with each location contributing one pixel to the output image.

Spatial filters come in a great many guises. The form and application of the most common filters are adequately covered in standard image processing texts (e.g. Castleman (1979), Niblack (1986), Rosenfeld and Kak (1976)). Filters may be classified under terms that reflect their mathematical properties. Terms such as linear, nonlinear, gated, and several others fit this bill. Alternatively, a common division into low pass and high pass filters is based upon the effect that the filters have on the frequency information contained in an image.

2.4.7 Low Pass (Smoothing) Filters

Low pass filters enhance low spatial frequency information, or features that are larger than the filter itself (Curran, 1985). They are often called 'smoothing' or 'defocus' filters due to their visual effect upon an image. The most common use for this class of filter is to enhance an image by the suppression of high frequency noise. The simplest smoothing filter takes the mean value of the filter window as its output. Davis et al. (1975) report using this process to suppress high frequency detail prior to thresholding to yield greatly improved results. The principal problem in low pass filtering is to avoid the blurring of objects boundaries while still reducing noise. Many of the nonlinear procedures, using both nonlinear filters and nonlinear implementations of linear filters, have proved most successful in this respect.

2.4.8 High Pass (Edge Enhancement) Filters

High pass filters enhance the high spatial frequency information, or features that are smaller than the filter itself (Curran, 1985). The most familiar type is the edge enhancement filter. Edge enhancement and edge detection are important stages in many automated image analysis applications. They are also a common form of enhancement prior to human interpretation. There are many designs of edge filtering available including the template matching and differential methods. Hall (1979) compares the merits of ten local edge detection procedures including the Roberts, Sobel, and Laplacian filters. Abdou and Pratt (1979) evaluate edge filter design by assessing the filters measurement of the edge magnitude, the probabilities of correct and false edge detection, and by employing a figure of merit calculation. They conclude that 3x3 differential detectors (e.g. Sobel and Prewitt) perform notably better than 2x2 differential operators (e.g. Roberts). Furthermore, the differential operators were found to be as good as the template matching detectors and required fewer calculations.

2.4.9 Rank and Range Filters

A class of nonlinear filters known of as rank filters have shown themselves to be particularly useful in image processing tasks. Rank filters take the pixel values from within the filtering window and place them into rank order. The output of the rank(i) filter is the ith value in the ranked sequence. If the extreme rank positions are selected the filters are known as **Min** and **Max** filters (Nakagawa and Rosenfeld, 1978). If the window contains an odd number of pixels the median rank produces a **median** filter. The properties of rank filters are

discussed by Hodgson et al. (1985). In particular the median filter is a powerful alternative to the low pass mean filter. It smooths noise but causes less degradation of edges. Rank filters can shift edges and alter the mean intensity of the image, but no new intensity values are generated. Nakagawa and Rosenfeld (1978) discuss the use of Min and Max filters for shrinking and expanding operations. These may be used to remove noise or detect densely clustered regions.

Range filters, described by Bailey and Hodgson (1985), are an extension of the rank filter. The output from a range filter is the difference in intensity between two selected positions in the ranked list of pixel values. Range filters may be used for edge detection and enhancement purposes, and by selecting an appropriate range the edge response width, edge position and connectivity scheme, may be adjusted. Bailey and Hodgson conclude that a range filter is more flexible than the Sobel filter but does not perform as well in a noisy image.

2.4.10 Efficient Filtering

Convolution is a term that is widely used when discussing spatial filtering. It consists of an operation whereby the output value at a point is formed from multiplying the image values by the filter coefficients and taking the sum total of these results. A substantial quantity of computation can be involved in this process if the window is large. Spatial filtering with windows larger than 11x11 pixels is generally less efficient than performing the convolution in the frequency domain using the Fast Fourier transform (Schowengerdt, 1983).

Calculating each value from scratch when the filter has only moved one

pixel position is often inefficient when symmetry exists in the filter coefficients (a common situation). Lee (1983) discusses the elimination of redundant operations in respect to producing a fast Sobel operator. Computation may be reduced by more than a half, the only penalty being increased storage requirements. The principles proposed by Lee are in fact applicable to most convolution filtering operations where a degree of symmetry in the filter design exists. Kim and Strintzis (1980) give details of an efficient convolution method which is especially effective as dimensionality increases.

Methods of efficient filtering with rank or median filters have also been investigated. Huang et al. (1979) employed a histogram storage technique. This algorithm is essentially equivalent to performing a 'running average' technique when using a mean filter. They demonstrated that the histogram method is much faster than an efficient sorting algorithm. Other methods using bit sorting (Ataman et al., 1980) and a sorted list (Bednar and Watt, 1984) are reviewed by Hodgson et al. (1985).

2.5 TEXTURE ANALYSIS

Texture is an important characteristic of many images. While tone and colour are obvious conveyors of information, texture also plays a considerable part. Variations in texture normally reflect some variation in the scene being imaged (Davis et al., 1981). This may be a change in the terrain, the type of skin tissue, the grain size of rock, or endless other possibilities. However, despite the fundamental importance of texture in image analysis it is still difficult to provide a simple definition of this phenomena.

Davis et al. (1981) provides one of the simplest definitions, "A textured area in an image is characterized by a nonuniform, or varying, spatial distribution of intensity". Others have emphasised the structural characteristic of many textures. For example, Wang et al. (1981) consider "...texture as composed of 'primitives' (connected regions satisfying certain properties) placed in a certain spatial arrangement". Tamura et al. (1978) give a similar description when they "...regard texture as what constitutes a macroscopic region. Its structure is simply attributed to the repetitive patterns in which elements or primitives are arranged according to a 'placement rule'". They go on to provide an equation,

$$\mathbf{f} = \mathbf{R(e)}$$

where a texture **f** is a function of a placement rule **R** and texture elements **e**. Irons and Petersen (1981) note the interrelation between tone and texture when they state "Visual texture refers to the impression of roughness or smoothness created by variation of tone or repetition of visual patterns across a surface". This theme was also discussed in some detail by Haralick (1979) who concluded that tone and texture are closely inter-linked. When there is little variation in tonal primitives, tone is the dominant feature, but as variation increases texture becomes the dominant partner.

From the definitions given above it can be seen that there is still debate and some uncertainty as to what precisely constitutes a texture. To some extent this uncertainty is reflected by the development of a wide range of techniques which attempt to quantify texture properties in digital images.

2.5.1 General Review

A comprehensive survey of techniques for texture analysis is provided by Haralick (1979). There are two approaches to texture quantification; statistical and structural. Statistical techniques attempt to characterize texture by statistical references to the spatial distribution of the grey levels. This is normally achieved on a pixel-to-pixel basis although local measurements are sometimes used. Structural analysis differs in that it aims to identify and define 'texture primitives', and then proceeds to specify their 'placement rules'.

Eight statistical approaches to the characterization of image texture are reviewed by Haralick. Autocorrelation functions, optical transforms, and digital transforms (including the Fourier transform), are similar in that they all analyse spatial frequency. This relates to texture in the sense that fine textures are rich in high spatial frequencies, whilst coarse textures are rich in low spatial frequencies. The Fourier transform provides directional as well as frequency information.

The 'textural edgeness' method considers texture as an amount of edge per unit area. This approach was employed by Rosenfeld and Thurston (1972) using the Roberts gradient as the local edge property. Sutton and Hall (1972) extended the technique to allow the identification of pulmonary disease by textural analysis.

Spatial distribution and dependence among grey tones in a local area may be analysed by the co-occurrence technique. Grey level run-length analysis characterizes texture in terms of the frequency spectrum of collinear connected sets of pixels that have equal grey tone. Both of these techniques are discussed in further detail later.

Texture transforms construct a new image which indicates how strongly a texture pattern is represented around each pixel of the original image. Since the textural properties at each pixel location are characterized, this approach is particularly attractive for remote sensing applications, where most classification algorithms assign categories on an individual pixel basis. Irons and Petersen (1981) attempted the thematic mapping of Landsat images with this method. They found the results to be poor but conclude that higher spatial resolution (now available with the latest generation of satellites and sensors) and better classifier design might lead to improvements. Furthermore, some transforms were found to produce useful image enhancements.

Haralick (1979) notes that pure structural based approaches have not been widely used and concentrates on statistical-structural hybrids that he divides into two classes. The 'weak' texture measures have only weak spatial interaction between primitives. Many of the statistical measures described above, such as run-lengths, edgeness, and local extrema density (Mitchell et al., 1977) fall into this class. The 'strong' texture measures also take into account the co-occurrence between texture primitives. The best example of this is the analysis of local maxima co-occurrence suggested by Davis et al. (1979).

Haralick concludes that for microtextures the statistical methods have been proven to work well, whilst for macrotextures the trend has been towards the use of statistical-structural generalizations based on histograms or co-occurrence of primitive properties. So far, structural methods have found success only in very limited applications.

2.5.2 Grey Level Run lengths

Analysis by grey level run length (Galloway, 1974) relies on recording the frequency of linearly connected sets of pixels having equal grey level. A fine texture is expected to have runs primarily of short length and very few of long length. Conversely, coarse textures will tend to have long runs of constant grey level occurring relatively often. A simple matrix may be used in which an element $\mathbf{P}_{(i,j)}$ represents the number of times that a run length of **j** pixels occur, having a grey level of **i**. Matrices may be computed for runs analysed in several directions, differences between these matrices reflecting properties of directionality.

This technique has been adapted for the analysis of handwriting using run lengths measured in the vertical and horizontal directions (Arazi, 1977). In similar experiments with ancient Hebraic handwriting (Dinstein and Shapira, 1982) over 91% of text samples were correctly classified, and 100% of individual letters. In these cases a binary image is used so the matrices condense into simple histograms. This method is particularly suited to the task since most characters consist of straight line segments which are readily characterized by their run lengths. In the analysis of terrain image samples the method was found to be susceptible to interference from noise (Weszka et al., 1976)

2.5.3 Grey Level Co-occurrence Analysis

Haralick et al. (1973) developed a set of texture measures for the classification or categorization of images based on a set of grey tone spatial dependence matrices. These matrices have since become

commonly known as grey level co-occurrence matrices. They specify the relative frequency with which any two grey levels occur in the image at some specified vector displacement. Fourteen measures of matrix distribution were proposed which could yield valuable textural information. Applying the technique to a selection of sandstone photomicrographs, aerial terrain photographs, and satellite imagery, good results were obtained with a piecewise linear discriminant classifier and independent test set.

It has been established for some time that important textural information is contained in second-order statistics (Julez, 1962). The co-occurrence matrices are a second-order equivalent of the histogram, specifying the probability of occurrence of any pair of grey levels rather than of a single grey level.

The power of the co-occurrence technique was confirmed by Weszka et al. (1976) in a comparative study with Fourier based descriptors and run length analysis. The co-occurrence technique produced decidedly better classification success rates. However, they also found that simplifying the matrices into vectors of grey level difference yielded results of a similar standard.

Rosenfeld et al. (1982) applied the grey level difference method for the simultaneous texture analysis of multispectral imagery. This type of imagery is common in remote sensing applications. Concentrating particularly on the two-band case they found that the texture features derived could provide information that was not available by single-band analysis.

Grey level co-occurrence analysis has been applied successfully to a wide

range of images and textures. Gersen and Rosenfeld (1975) employed it to distinguish between clouds and sea ice on aerial photographs. Photomicrographs of various sandstones have been differentiated on the basis of their co-occurrence texture descriptors (Haralick and Shanmugam, 1973). Bertolini and Vernazza (1982) utilized co-occurrence matrices for the analysis of nucleated cells, and Don et al. (1984) found it possible to assess the roughness of metal surfaces using co-occurrence statistics. Chien and Fu (1974) applied grey tone co-occurrence to automated chest X-ray analysis.

Co-occurrence analysis has frequently been used to provide texture measures for image segmentation. Chen and Pavlidis (1979) employed co-occurrence matrices with a split-and-merge algorithm, while Mason (1979) used the simpler grey level difference vectors of Weszka et al. (1976). Sklansky (1978), while providing a general overview of image segmentation and feature extraction processes, notes the use of co-occurrence as a common and important texture analyzer.

2.5.4 Generalized Co-occurrence

Davis et al. (1979) developed co-occurrence statistics further with their introduction of generalized co-occurrence matrices (GCMs). The GCM approach, one of Haralicks (1979) 'strong' texture measures, utilizes co-occurrence of local property values that satisfy certain spatial constraint predicates. For instance, the local properties of edge strength and orientation could be obtained. Possible spatial constraint predicates might be that the edge properties are nearest neighbours or that they occur within a threshold distance of each other.

This method is particularly aimed at quantifying coarser or macro textures which are less appropriately analysed by standard grey level co-occurrence. As texture becomes coarser the statistics derived from normal grey level co-occurrence will increasingly reflect the intensity transitions within the texture elements rather than the structural organization of the texture (Shen, 1980).

Davis et al. (1979) discovered that with a small data set of 30 samples from five texture classes, and using a leave-one-out euclidean distance classifier, the GCMs yielded over 80% accuracy compared to 50-57% for standard grey level co-occurrence analysis. In a later study a much larger data set of 128 images from eight texture classes were used (Davis et al., 1981). Texture descriptors were evaluated from grey level co-occurrence matrices, edge-pixel GCMs, and extended-edge GCMs. Entered into three classification schemes the grey level co-occurrence statistics gave slightly better overall results. This was attributed to their superior performance when comparing textures of similar characteristics within the data set available. When dissimilar textures were compared the GCM measures provided the greatest levels of statistical separation.

Dyer et al. (1980) attempted to combine the grey level co-occurrence and generalized co-occurrence techniques. They proposed a scheme in which grey level co-occurrence matrices are constructed using pixels at locations and separations defined relative to the position and orientation of local edge maxima. In this way account is taken of both the tonal and structural properties of a texture. Furthermore, the statistics should be less sensitive to the coarseness or size of primitive elements in the texture. Improved results over the standard grey level technique were reported for a small experimental data set.

2.5.5 Local Property Analysis

Pietikainen et al. (1983) developed a statistical scheme for texture analysis based on the analysis of local image properties. The texture features describe the degree of match between each pixel neighbourhood and a set of standard masks. Masks of both 3x3 and 5x5 dimensions were convolved with the image and the sum of absolute values employed as the texture descriptors. The design of the masks made them sensitive to particular local properties such as edges, lines, spots, and ripples. Comparing results on a small data set of 28 images they concluded that the local property masks had greater discriminating power than statistics based on grey level co-occurrence.

2.5.6 Structural Analysis

Wang et al. (1981) describe three schemes for the explicit extraction of texture primitives. The simplest uses fixed percentile thresholding of the grey level histogram to extract the upper or lower 25 percent of points in the image. The second used the histogram modification scheme of Peleg (1978) to condense the histogram into a few spikes, containing around 25 percent of the image points. Finally a 'superslice' algorithm was used which segments the image by a selection of thresholds, where the threshold level generating the greatest coincidence with an edge map is selected. This last technique allows information from both the histogram and local properties to be taken into account.

Once primitives had been extracted both first-order and second-order statistics were calculated. The first-order statistics measured size, shape and direction properties of the primitives. The second-order statistics

were derived from primitive attribute co-occurrence matrices. These were employed in order to capture some information on the spatial distribution or 'placement rules' of the primitives. Applying these methods to a data set showed that the superslice extraction method gave noticeably better results with second-order statistics than the others. Also, the second-order statistics could distinguish between some textures where the first-order statistics failed. The overall conclusion, however, was that the simplest methods had worked about as well as the more complex.

Hong et al. (1980) discuss a structural technique in which the extraction of primitives was achieved by the use of edge pairs rather than the more common threshold based schemes. Eight 3x3 gradient template detectors were used to detect edge points. The edge map was then thresholded to remove weak edges and non-maxima values suppressed. Edges were paired if they had opposite orientation and were positioned within a maximum threshold distance from each other. An interior point probability value was increased whenever an image point occurred between paired edges. Finally, a weak region growing stage was used to clean up the extracted primitives. The mean values of area, perimeter, dispersiveness, and other attributes, were evaluated from the extracted primitives to provide the texture measures.

Pietikainen and Rosenfeld (1982) produced a similar class of texture measures based on the first-order statistics derived from edges in an image. The same scheme as Hong et al. (1980) was used to obtain an edge map. Edge pairs of opposite orientation and within a threshold distance were identified in the belief that these would relate to the texture primitives. The mean and range of grey levels of points between pairs, the separation of the edge pair, and the difference in

edge contrasts were calculated. These were used along with the mean and variance of edge curvatures as texture measures. This statistical-structural hybrid is closely related to the generalized co-occurrence matrices of Davis et al. (1979). The texture measures yielded good discrimination amongst scenes of different geological terrain. This method does not require the explicit extraction of primitives, and the texture measures have the advantage of simple perceptual interpretation.

2.6 ANALYSIS OF SHAPE AND GEOMETRIC PROPERTIES

Considerable research has gone into analysing the geometric properties of objects captured in digital images. Much of this work has been aimed at robot vision systems to allow the discrimination, sorting, and manipulation of parts for automated robotic assembly. In general these parts are of well defined and consistent form. Therefore, there is considerable advantage to be gained in the analysis of both simple properties such as area, compactness, and centre of gravity, as well as more complex properties such as principle axes, the number, magnitude and location of corners, and so on. Umetani and Taguchi (1979) define a number of properties suitable for discriminating between complex shaped machine parts.

Pollen viewed under an optical microscope has reasonably consistent shape and is an important feature used by palynologists in manual identification. However, under the SEM pollen is frequently distorted and is no longer a reliable feature for classification purposes. Therefore in an automated system of pollen identification it is unlikely that there will be any great profit in the detailed analysis of shape.

Simple properties such as area and compactness may help distinguish between pollen grains and other artifacts in the image. Furthermore, most pollen have a degree of symmetry and corners or points on their outline, whilst not precisely located, are likely to be evenly spaced and of roughly similar magnitude. Therefore the calculation of simple geometric properties of pollen in the digitized image may provide valuable information for discriminating between pollen grains and other unwanted debris.

2.6.1 Boundary Encoding

Many geometric properties of an object can be obtained from the analysis of its boundary. Freeman (1961) proposed a simple yet compact method of encoding such boundaries known as chain code. Batchelor and Marlow (1980) describe a hardware/software mechanism which allows rapid generation of chain code from binary images. Dudani (1976) provides an algorithm for the location and following of object boundaries and the generation of a binary image displaying the extracted regions. Other boundary following techniques are given by Rosenfeld and Kak (1976), and by Bennet and Mac Donald (1975).

2.6.2 Geometric Properties

Freeman (1961) illustrated that it is possible to determine boundary length, and enclosed area from a chain coded description of the object. The calculation for centre of gravity is a simple extension of the area calculation, and Kakikura (1979) provides a fast method of extracting the feature axes of objects from their boundary chain code.

A classic measure of compactness for a figure is the square of the perimeter divided by the area. Rosenfeld (1974a) notes that this measure is less satisfactory with digital images since the smallest values may be yielded from octagons or diamonds rather than circles. Since these precisely defined forms are most unlikely to occur with pollen images this problem should not be important. However, Haralick (1974) provides a simple measure for the circularity of digital figures which behaves well.

As stated earlier the precise location and details of corners are unlikely to be of great benefit for pollen identification. However, since many pollen taxa have an approximately smooth and symmetrical outline it may be possible to differentiate between pollen and other artifacts using shape information.

The principle problem in detecting corners from chain coded objects is that the spatial quantization in a digital image causes most straight lines to be represented by a series of alternating slope segments. These characteristics are examined in some detail by Rosenfeld and Kak (1976), and by Rosenfeld (1974b) who defines a test for a true digital straight line.

Davis (1977) describes several methods that have been proposed for angle detection and provides a technique for hierarchical shape analysis. The simpler method proposed by Rosenfeld and Johnston (1973) is more likely to be suitable for analysing pollen images. This calculates rates of curvature by detecting the differences between successive non-overlapping slope segments, incorporating a variable degree of smoothing. An improved version of this method is presented by Rosenfeld and Weszka (1975).

Freeman and Davis (1977) also employ a 'smoothing line' connecting two points on the chain coded boundary, but claim their method is simpler and requires less computation. A 'corner prominence' measure is also described, based on the magnitude of the angle and the length of the straight line segments on each side. They illustrate the effectiveness of their method with several examples.

CHAPTER THREE

METHODS AND EQUIPMENT

3.1 INTRODUCTION

This Chapter describes the source and the type of pollen that were used in the experiments on automated identification procedures. It also reviews the equipment that was used to obtain suitable images of these pollen and describes the computer hardware used to digitize these images. Finally, both the hardware and the software used to analyse the image data is discussed.

3.2 POLLEN TAXA USED IN THE EXPERIMENTS

Only fresh pollen were used during the course of these investigations. It was felt that fossilized pollen would add unnecessary complications at this early stage of experimentation. Fossilized pollen samples extracted from sediments will contain unwanted materials no matter how thorough the preparation techniques used. The unwanted material normally consists of amorphous organic matter and silica spicules. This clutter hinders the locating of pollen under the microscope. Fossilized grains are also likely to have suffered some degree of distortion. In very severe cases the grains can be flattened and fragmented making the identification process many times more difficult, even for the human observer.

Six pollen taxa were utilised in this work, these are shown in Table 3.1 below. For simplicity, the report refers to these taxa by their common rather than taxonomical names from now on. An example of a typical pollen grain of each taxon is illustrated in Plate 3-I.

Code	Taxonomical name	Common name
(Pi)	Pinus sylvestris	Pine
(El)	Ulmus glabra	Elm
(Oa)	Quercus robur	Oak
(Ry)	Lolium perenne ssp. multiforum	Rye grass
(Ha)	Corylus avellana	Hazel
(Pl)	Plantago lanceolata	Plantain

Table 3.1 Pollen taxa used throughout these experiments.

3.3 THE PREPARATION AND VIEWING OF POLLEN SAMPLES

Since the fresh pollen was not encapsulated in any sedimentary material, most of the normal preparation techniques used to remove the extraneous substances were not needed. The process of acetolysis was employed to strip out the inner cellulose layers so that the pollen more closely resembled a natural sediment derived sample. The samples were dried at room temperature (since equipment for critical point drying was not available) and mounted on aluminium stubs before sputter coating with pure gold.

The prepared samples were viewed on a Cambridge 600 Series Stereoscan scanning electron microscope. Photographic images were taken using a built-in camera and phosphor screen, once suitable grains had been located. The negatives were processed to provide black and white prints of about 15x20cm dimensions. The SEM was selected in preference to an optical microscope as it could provide images with much greater detail. Furthermore, a complete pollen grain could be kept in focus due to the far superior depth of field available.

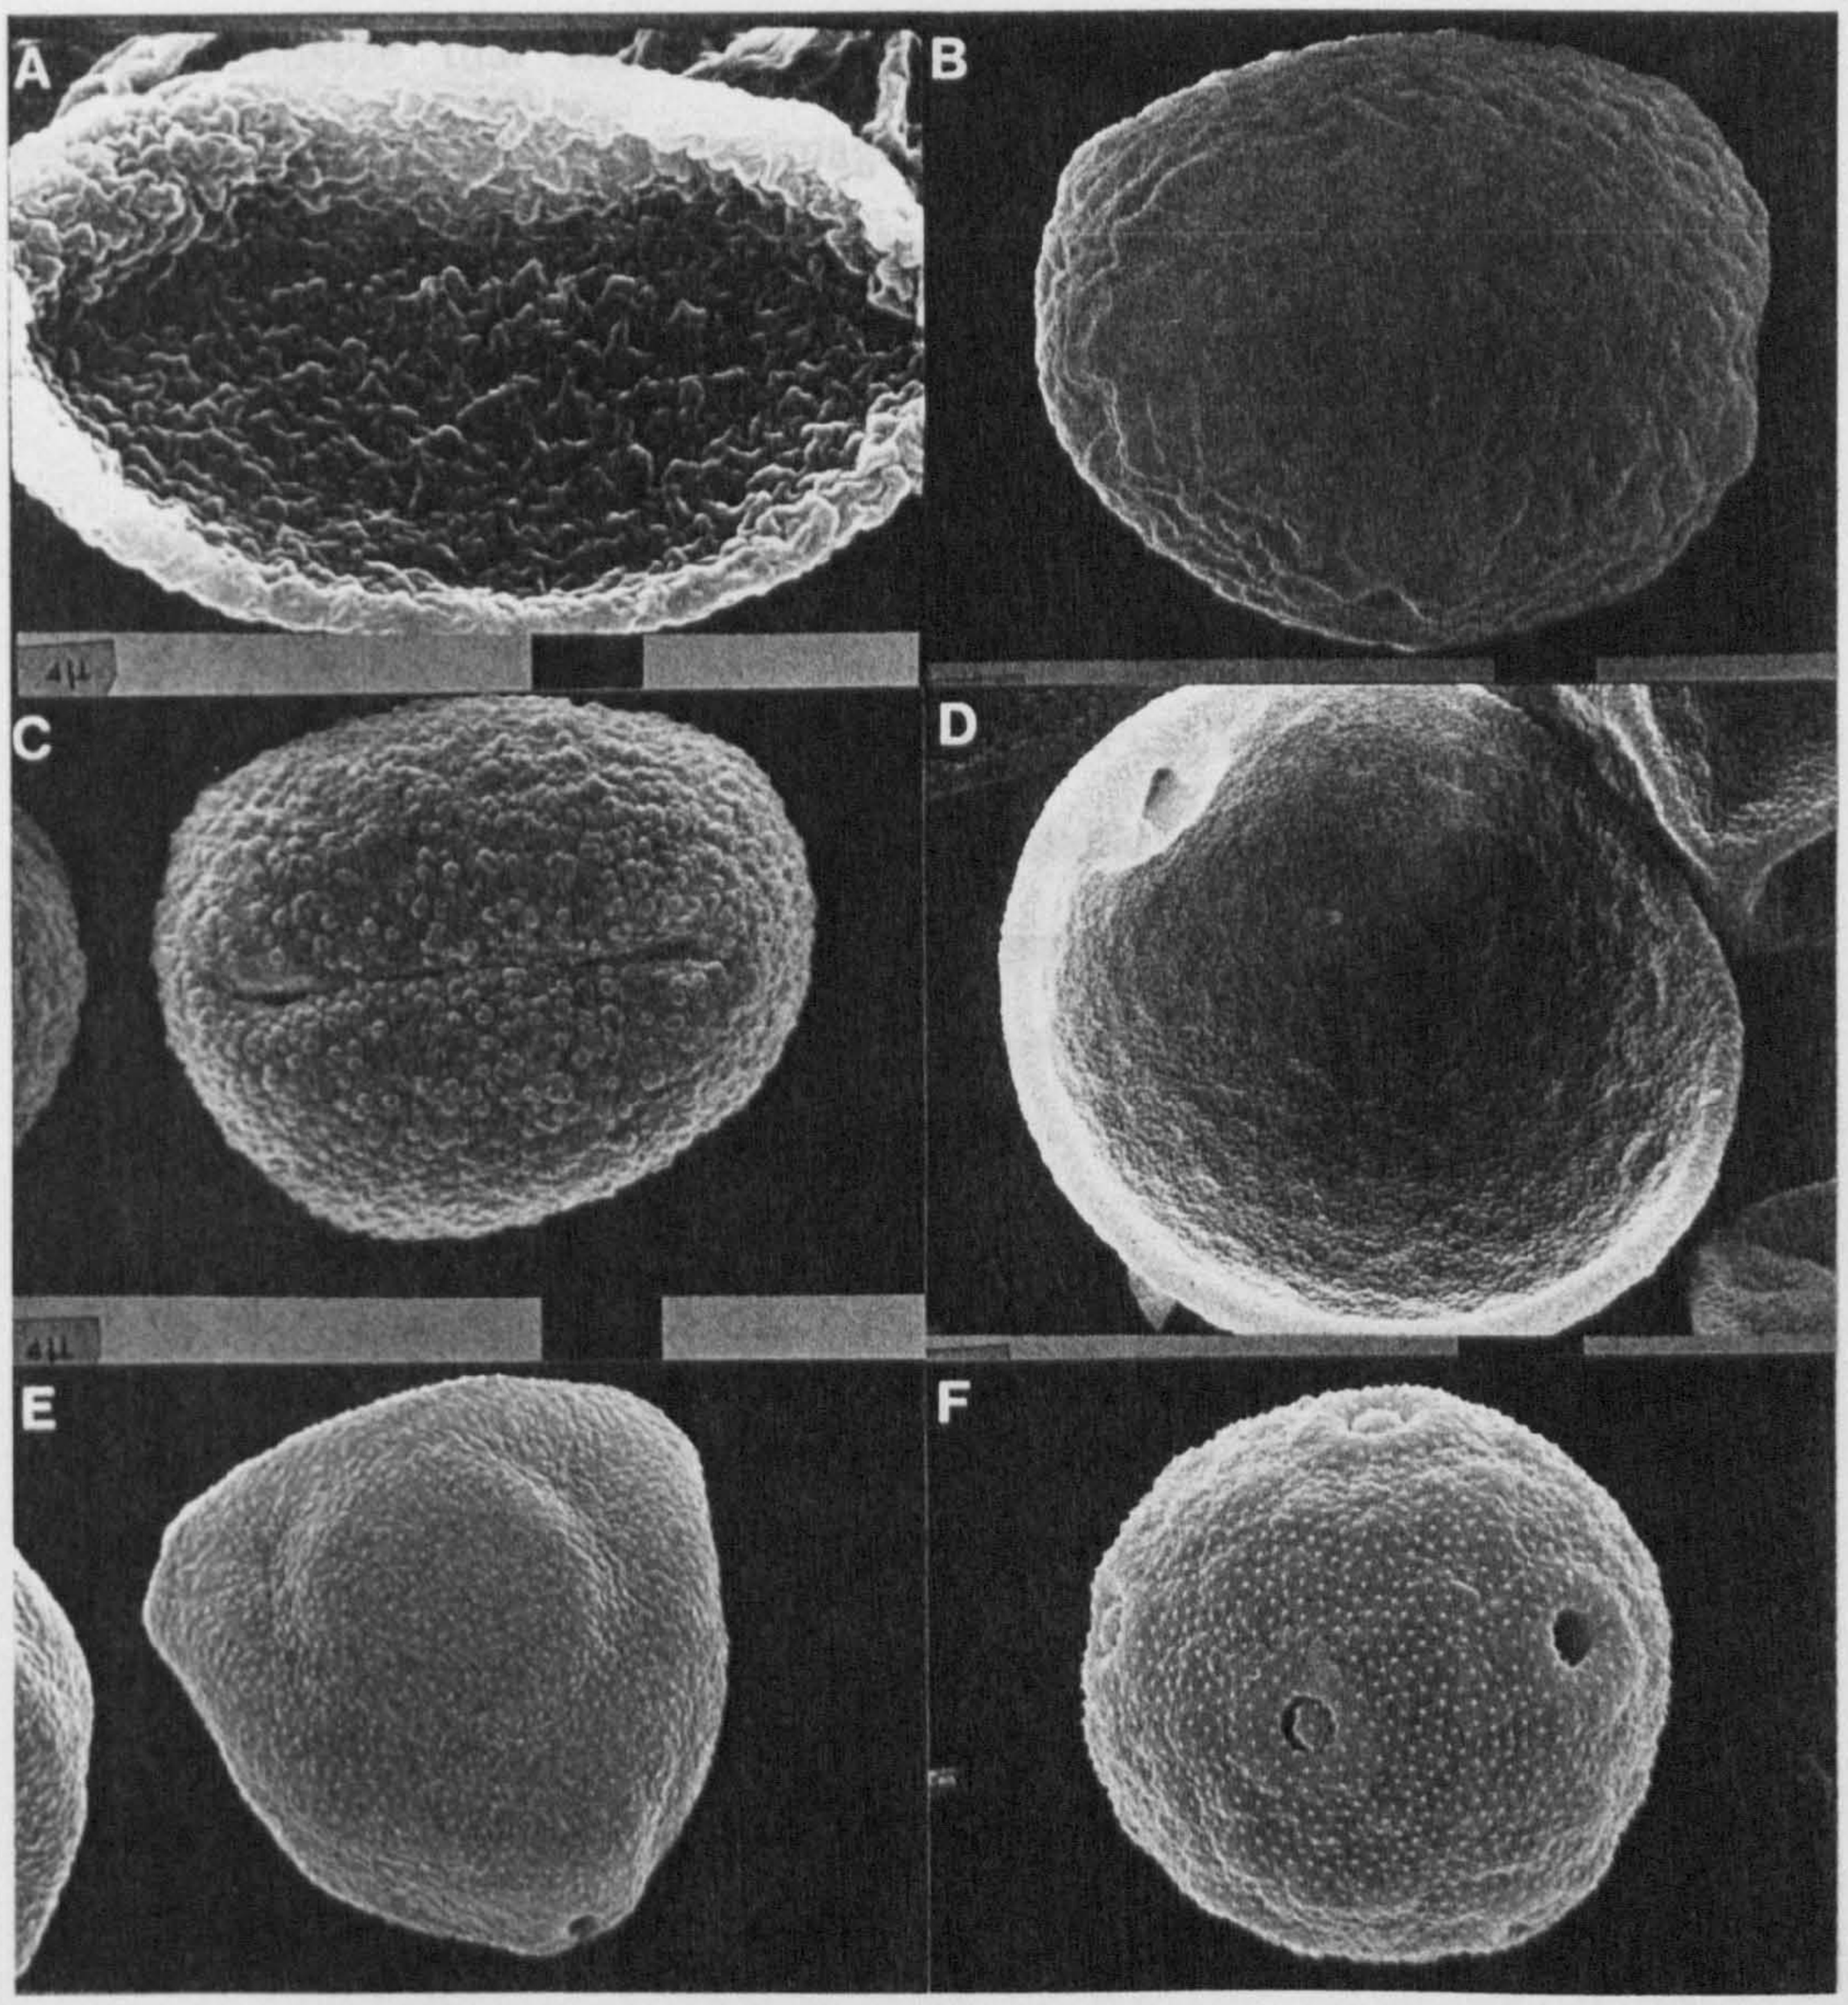

Plate 3-I. An example of each pollen taxon, clearly showing the detail and depth of field that is obtained on the Scanning Electron Microscope (SEM). (a) Pine, (b) Elm, (c) Oak, (d) Rye grass, (e) Hazel, (f) Plantain.

The magnification on the SEM was selected to try and ensure that a complete pollen grain would fit comfortably inside the screen boundaries. Unfortunately, the SEM in use had a fairly coarse magnification adjustment so the best overall setting could not view a complete Pine grain since these are larger than the other taxa. The final prints produced had a magnification of approximately X2000. Efforts were taken to ensure that the contrast and brightness levels were kept as constant as possible for each image captured, but this was very difficult to achieve especially between viewing sessions.

3.4 IMAGE PROCESSING AND CLASSIFICATION SOFTWARE

Software was written almost exclusively in FORTRAN. This was mainly due to the greater level of support available for this language. The presence of previous in-house image processing software written in FORTRAN was a further incentive, although this was only used during the earlier days.

Other in-house utility routines were available on the image digitiser. These consisted of simple pixel read and write facilities on the frame store, along with image capturing, saving and loading procedures. These were used quite extensively throughout the project. In particular the image capture and save routines were used to convert the photographic images to a digital format.

3.5 IMAGE PROCESSING HARDWARE

All digital images of pollen were derived from the SEM photomicrographs. They were captured and digitized using a Cotron video camera attached to a Matrox frame store. The frame store had a 256x256 pixel resolution and eight bit planes providing 256 intensity levels. The system was controlled by a DEC PDP 11/23 minicomputer running the RTS single user operating system. Plate 3-II illustrates the digitization equipment in use.

In addition to capturing and storing images, the PDP 11/23 could be used to perform some basic image processing tasks. Generating grey level histograms to check the dynamic range utilised during image capture, and contrast manipulation techniques, are examples of the procedures that could be undertaken. It was not uncommon to test out ideas on this machine before moving the images to a more powerful computer for serious processing.

Much of the early image analysis work was performed on a DEC PDP 11/34 machine running the RSX multi-user operating system. It was possible to view the images using a frame store attached via a parallel interface.

This machine was not entirely suitable to the task of image processing for a number of reasons. The most important of these was the very limited memory available which made it possible to store only a few lines of the image at any one time. A considerable amount of input/output operations were therefore required for most processing tasks. Naturally this severely limited the speed of processing and necessitated efficient rather than elegant coding. With the multi-user environment

Plate 3-II The equipment used to digitize pollen images. This consisted of a video camera for viewing the photomicrographs, attached to a Matrox frame store and display device. These were all under the control of a DEC PDP 11/23 computer.

the problem was further exacerbated when the machine was placed under heavy load by other workers. A lesser problem was the limitations imposed by a FORTRAN IV compiler.

After twelve months of work the PDP 11/34 was replaced by a DEC VAX 11/750 minicomputer operating under VMS. This provided much greater processing and data handling power. It was possible to store several complete image arrays in memory at a time. A FORTRAN 77 compiler was available on this machine.

There were, however, two disadvantages with the VAX machine. Firstly, a file transfer package had to be used to transfer images from the PDP 11/23 system where they were digitized. This was a fairly slow procedure. Secondly, it was not possible to view the images once on the VAX as there was no frame store attached. Operations that required viewing the image had to be carried out on the PDP 11/23. Fortunately these all occurred at the early stages of processing, after which the image could be transferred to the VAX for serious 'number crunching'.

CHAPTER FOUR

IMAGE PROCESSING TECHNIQUES

4.1 APPROACHES TO THE IDENTIFICATION PROBLEM

4.1.1 Finding and Identifying Pollen

Before discussing the image processing techniques that were used to help identify pollen we need to consider the problem in a little more detail. For an automated pollen identification system there are in fact two major problems to solve. The first is the need to 'find' pollen under the SEM, and the second is to 'identify' these objects.

The initial task requires the computer to control the microscope stage and to have the ability to locate objects when they appear within an image. The sort of problems to be solved are,

(a) Are there any objects in the image?

(b) If so, how many, and where are they?

(c) For each object, is it fully visible or will the microscope stage need to be moved a little in order to see the entire object?

(d) How likely is it that the object is a pollen grain and not some other artifact?

These problems are not dealt with in detail in this report; time did not allow a thorough investigation into them. In Chapter Seven a brief review of some techniques that might prove useful are presented, but essentially the task remains an area for future research. Instead, the problem of identifying pollen is to be tackled here, making the assumption that some scheme for providing the location of the objects will be developed in the future.

4.1.2 Potential Features for Identification

Given the problem of identifying a pollen grain a decision has to be made as to what features are to be measured in order to perform the task. Potential candidates are the size and the shape of a grain, the number and location of apertures in the exine surface, and the pattern of exine sculpturing or texture.

When pollen is identified manually under an optical microscope it is the number and location of apertures, and the size and the shape of a grain, that are the primary sources of discriminatory evidence. Due to the low resolution of optical microscopy surface texture is of secondary importance. However, the human's power of interpretation is a complex subject, and the precise importance of any feature is impossible to state as they all tend to be analysed simultaneously in the decision making processes of the brain. After consultation with experienced palynologists the following conclusions were made on the potential use of the stated features for automated identification of pollen using SEM derived imagery.

4.1.3 Grain Shape

Grain shape was considered to be the least reliable property. Due to the severity of the SEM environment fresh pollen frequently become distorted. Fossilized pollen is typically squashed or distorted before it is even subjected to SEM analysis. There is an additional complication in that the SEM is able to view objects obliquely. Unless this was carefully avoided, even a well formed grain would be difficult to identify due to the distortion of shape imposed by the viewing angle.

4.1.4 Grain Size

Grain size was considered to be a slightly more reliable feature. However, since pollen is biological material its size does vary even for objects from the same pollen taxon. This imposes limitations on its reliability and diagnostic power. Almost all pollen fall within the size range of 1.5×10^{-7}m to 5×10^{-7}m. Therefore, this property might perhaps be more usefully employed in answering question (d) posed in Section 4.1.1 above, i.e. how likely is it that the object is really a pollen grain?

4.1.5 The Number and Location of Apertures

Apertures within the exine are an extremely important feature in manual identification using optical microscopy. Their number and location are usually the first feature analysed by the palynologist. It must be appreciated, however, that it is possible to view the complete pollen grain when using optical microscopy. This may be achieved either by focussing at successive levels though the grain (one advantage of the narrow depth of field), or by rotating the grain (which is possible since a fluid mountant is used). Consequently all the apertures may be seen. The image produced by a SEM does not allow this facility, only a fixed view of a solid object is provided in any single scene. It is uncertain how many pores are hidden on the unseen side of the grain, which severely limits the diagnostic power of an aperture count.

4.1.6 Exine Texture

Exine sculpturing, or surface texture, was believed to be the most reliable and power feature for computer based discrimination of pollen taxa. Almost all pollen taxa have raised and depressed areas on the exine giving rise to a surface texture. Under the SEM these are shown in great detail and the texture proves to be highly characteristic of a given pollen class. Although these patterns do vary slightly due to the nature of biological materials, they appear to be reasonably consistent. Moreover, since they cover a large area of the grain we might assume that statistical properties of the texture can overcome the problem of small scale variability, providing a sufficiently large region is analysed.

4.1.7 Conclusion

From the discussion presented above it should be self-evident that exine texture is the most promising feature for automatically discriminating pollen within the SEM environment. Section 4.3 describes the methods that were used in an attempt to measure the textural properties of the exine surface from the six pollen taxa considered in this work. However, before this, Section 4.2 reports on the other digital image processing techniques that were necessary in order to allow the texture analysis schemes to proceed.

4.2 DATABASE PREPARATION TECHNIQUES

4.2.1 Extracting Sub-scenes of Exine Texture

Texture is a spatially based phenomena; it is a property that exists within an area and not at a point. Consequently it must be measured over an area, although how big this region should be is a topic for debate. In these experiments sub-scenes of 64x64 or 60x60 pixels were used for the calculation of textural properties. The actual area of the pollen grains covered by samples of this size obviously depended on their magnification in the digital images. In general approximately one tenth of the pollen surface area visible in an SEM image was covered by a sub-scene sample.

The location of exine texture regions was performed manually. Ultimately an automated method would be needed, but this should be a much simpler task than locating the grains themselves. While selecting the samples some care was taken to ensure that they were a representative set. Obviously apertures were avoided, but some samples were located close to the edge of the grain where the texture becomes distorted by the curvature of the surface. Moreover, the photomicrographs of pollen were positioned randomly under the video camera to ensure that there was no favoured orientation of textures displaying directional properties.

Since at most a 64x64 pixel region was to be used to measure the textural properties of a grain, storing the full image would have been a inefficient use of resources. Each 256x256 pixel image requires 64K of storage space, whereas the 64x64 pixel sub-scenes require only 4K. Therefore, the texture samples were cut out from the original full size

scenes, and recombined in a montage image file. In this way sixteen samples could be stored as a single image file. The effect of this process is clearly seen in Plates 6-I to 6-XII of Chapter Six.

4.2.2 Point Operations

Point operations are a relatively simple collection of procedures that transform an image on a pixel-by-pixel basis without regard for the values of surrounding pixels. The mapping between the input and output pixel values is defined by a transformation function. If we assume that an image has grey levels denoted by g. The transformation function has the form,

$$s = T(g)$$

where g is an original grey level, s is the new grey level it is mapped to, and T is the transformation function which normally satisfies the conditions of being single valued, non negative, and monotonically increasing.

Point operations are typically implemented by a table look-up procedure that allows rapid and efficient processing, especially on modern display devices using hardware directed look-up tables. Their main applications are in the manipulation of contrast, pseudo-colouring and thresholding.

4.2.3 Histogram Equalization

Digital images are frequently captured using only a small portion of the full dynamic range of intensity levels. When displayed, these images have a low level of contrast making it difficult to discern detail. A

popular enhancement is the linear contrast stretch which scales the intensity values to fill the full range available. This transform, and others of a more elaborate design, are invaluable for enhancing contrast levels to allow easier human interpretation.

Histogram equalization is a point operation that is commonly used to enhance contrast for human observers. However, it also has considerable value for automated image analysis systems. The transformation function used for histogram equalization is in fact the cumulative distribution function of the original image. The grey level probability density function of a digital image is represented by the equation,

$$P(g) = \sum_{i} \sum_{j} F(i,j) = g$$

where P(g) is the probability of grey level g occurring in the image F. The cumulative distribution function (CDF) is given by,

$$CDF(w) = \sum_{1}^{w} P(w)$$

It can be easily demonstrated that using the CDF as a transformation function should yield an image whose histogram has an even distribution among all grey level bins (Gonzalez and Wintz, (1977), Niblack (1986)), hence the term histogram equalization. However, only an approximation to a flat histogram can be achieved with a digital image due to the discrete nature of the data.

Using the CDF as a transformation function distributes contrast evenly across the full dynamic range and is frequently used as an enhancement

process. However, more importantly for automated image analysis, histogram equalization transforms an image into a standard format. The first-order statistics of the transformed image, in particular the mean and the variance of the intensity levels (which equate to the properties of global brightness and contrast) are standardized. This allows the direct comparison of images taken at different locations in time and space, since effects such as unequal lighting conditions at the time of image acquisition are effectively removed.

The global brightness and contrast of the digitized pollen images could not be tightly controlled. These parameters could be altered by the controls on the SEM, during the developing of the photographs, and during the digitization process. Many measurements obtained with the texture analysis routines described later were sensitive to these global parameters. Therefore, histogram equalization was commonly employed to normalize the first-order statistics of the pollen sub-scenes before texture analysis was undertaken.

4.2.4 Local Operations

Local operations transform an image on a pixel-by-pixel basis. However, unlike point operations, the value of a pixel in the output image is a function of the pixel values within a local area of the corresponding pixel in the input image. The application of local operators, or spatial filters as they are frequently known, is very common for image restoration, enhancement, and analysis purposes.

The term convolution is often employed when discussing spatial filtering. This consists of a repetitive 'shift-multiply-add' procedure which is

illustrated in Figure 4.1. The local neighbourhood is defined by a 'filtering window' which is typically square and symmetrical around a central pixel location. This window is placed exhaustively at all locations within the image area. At each location the intensity values of the underlying image pixels are multiplied by the filter coefficients. The resulting values are summed and this produces the new value that is assigned to the corresponding pixel in the output image.

The choice of filter coefficients determines what effect the operation has upon the image. Figure 4.1 illustrates the coefficients that are used for the operations of 'mean smoothing' and 'Laplacian edge detection' using a 3x3 pixel filtering window. In Chapter Seven a mean smoothing filter is used before thresholding to reduce high frequency information generated by texture. The convolution algorithm is an integral part of the Laws' mask texture measures described in Section 4.4, and it is also used for edge detection by the edge pair texture measurement scheme described in Section 4.5.

The nonlinear median filter used for noise reduction in the edge pair texture measurement scheme works on a slightly different basis. In this case the image pixel values within the filtering window are sorted into rank order rather than multiplied by coefficients. The median value in this ordered list is then used as the output value.

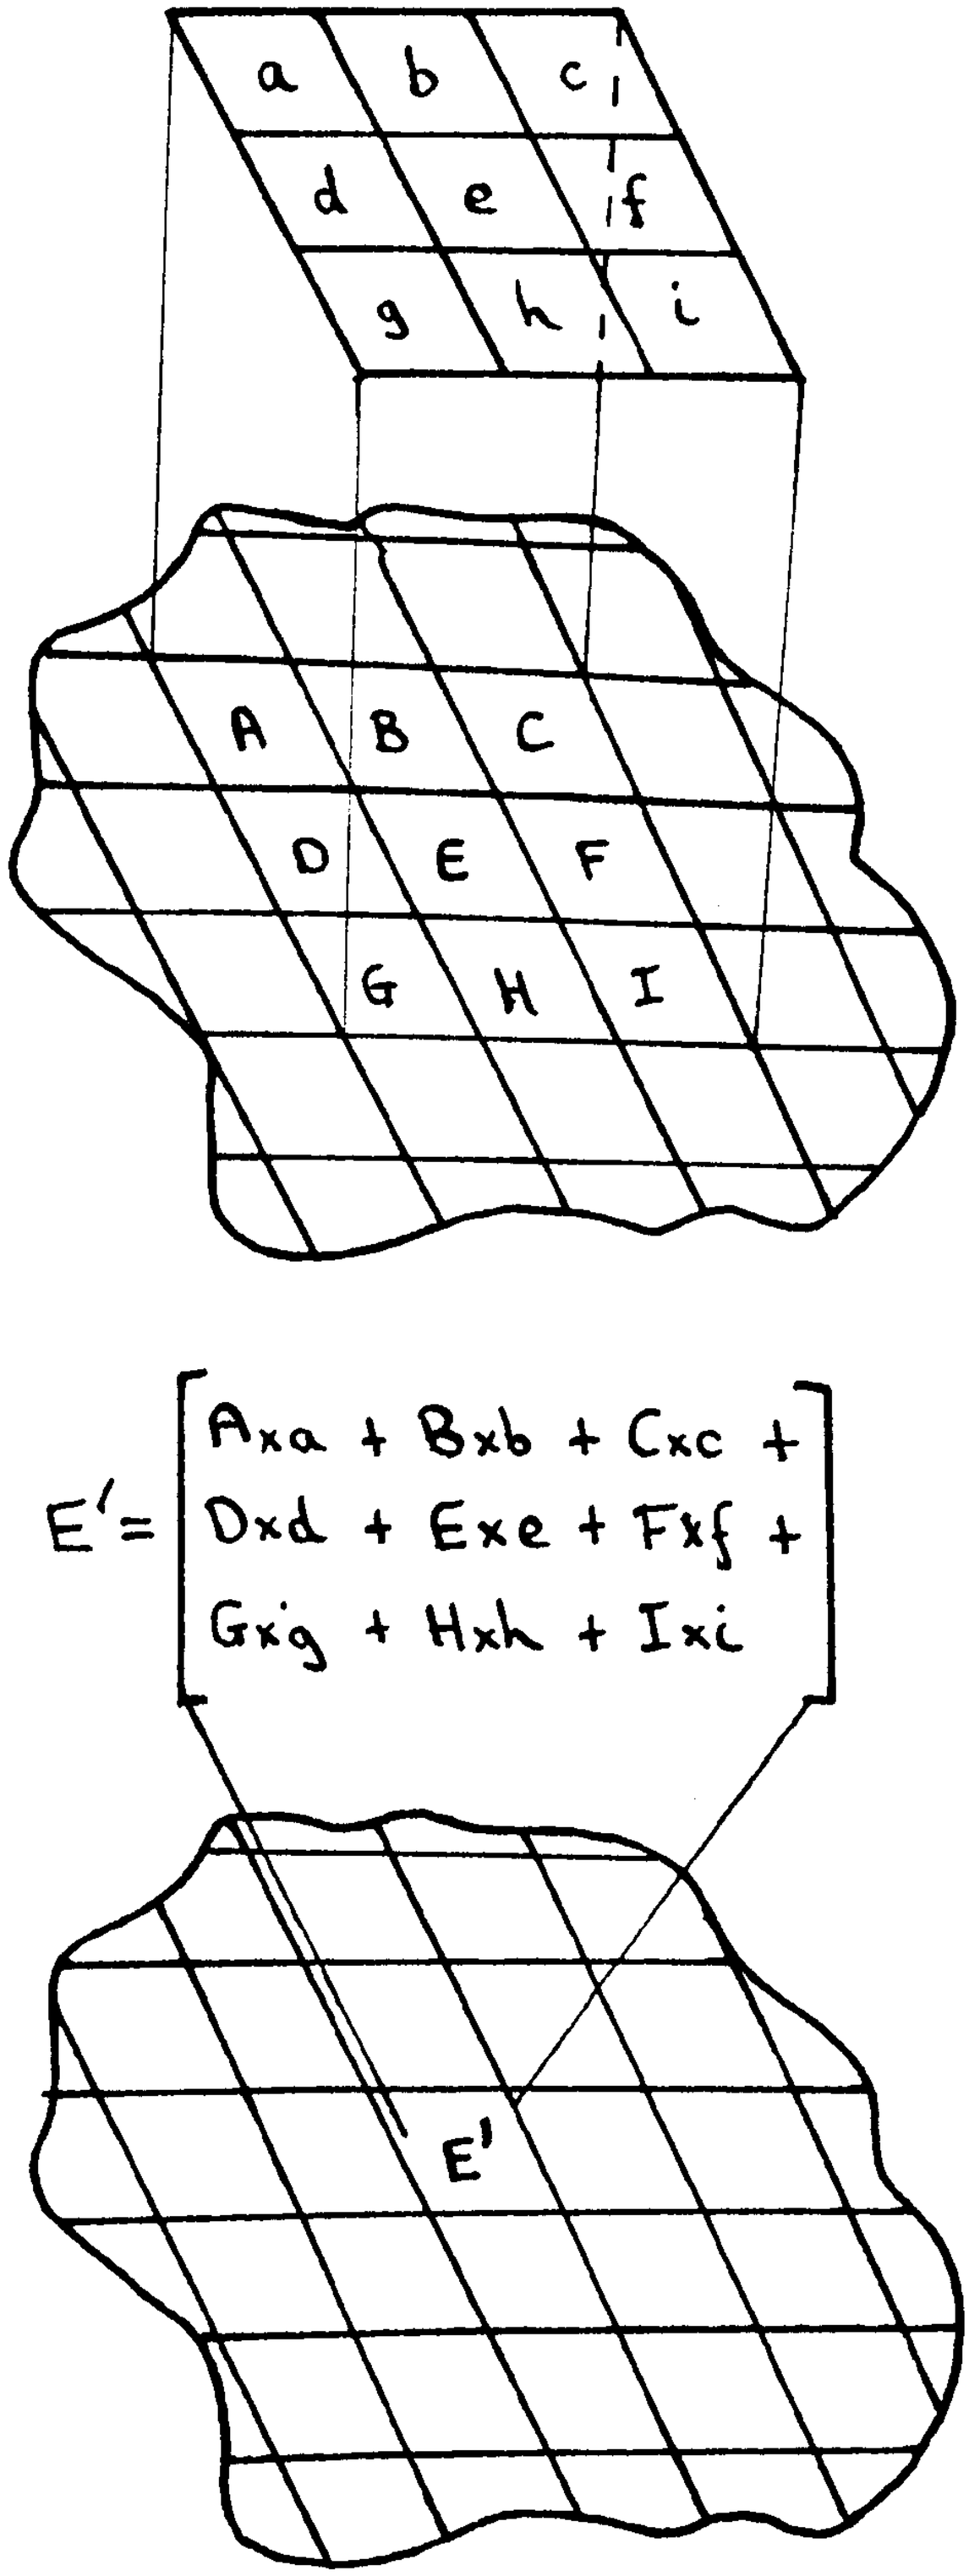

$\frac{1}{9}$ x

+1	+1	+1
+1	+1	+1
+1	+1	+1

Mean smoothing filter

-1	-1	-1
-1	+8	-1
-1	-1	-1

Laplacian filter

Figure 4.1 Spatial filtering by convolution with a local operator

TEXTURE ANALYSIS SCHEMES

4.3 GREY LEVEL CO-OCCURRENCE

Grey tone spatial dependence analysis has been one of the most successful and widely used techniques for the quantification of textural properties in digital images. It makes use of a grey tone spatial dependence matrix which has become commonly known as a co-occurrence matrix. Haralick et al. (1973) were the first to propose that statistics extracted from a co-occurrence matrix could be used as measures of textural properties. Since its introduction the technique has been used by many workers on a broad range of image types. Some examples that illustrate the diverse range of applications found for co-occurrence analysis are given below.

4.3.1 Examples Of Grey Level Co-occurrence Analysis

Haralick et al. (1973) introduced the technique by applying the measures to a classification problem involving three types of image data. These were photomicrographs of sandstone taken from five classes, aerial photography illustrating eight land-use categories, and multispectral satellite imagery containing seven land-use categories. Classification results of over 80% for all three image groups indicated the potential versatility and power of the extracted features.

Widely contrasting applications were found by Bertolini and Vernazza (1982), who employed co-occurrence measures for the identification of changes in biological cells, and by Don et al. (1984) who successfully quantified the roughness of metal surfaces. Gersen and Rosenfeld

(1975) used these descriptors in their attempts to differentiate between clouds and sea ice on aerial photography, while Chien and Fu (1974) found a biomedical application in the automated analysis of chest X-ray images.

The power of co-occurrence measures in comparison with other texture analysis techniques, including those based on the Fourier power spectrum and grey level run-lengths, was assessed by Weszka et al. (1976). They concluded that the texture features extracted from co-occurrence matrices. and the closely related grey level difference vectors, performed better than either the Fourier descriptors or the run-length measures.

4.3.2 Constructing Co-occurrence Matrices

The co-occurrence matrix can be considered as an extension of the grey level histogram to a higher-order distribution (Sklansky, 1978). The histogram simply represents the probability distribution of grey levels measured from individual pixels. The co-occurrence matrix represents the probability distribution of grey level pairs, where the two pixels providing the grey level values are separated by some pre-specified displacement vector, $d=(\delta x, \delta y)$.

Any element of the co-occurrence matrix $\mathbf{P}_{(i,j)}$ represents an estimate of the probability of finding the grey levels i and j at the spatial displacement, or relative positions, defined by the displacement vector. The dimensions of the matrix will be N^2, where N is the number of discrete grey levels present in the digital image. A simple example will help to explain the construction of a matrix.

Consider the very simple digital image depicted in Figure 4.2. This image has a spatial resolution of 4x4 pixels and a tonal resolution of four grey levels, ranging in value from 0 to 3. The first step in constructing a co-occurrence matrix is to define a 'mask' which specifies the displacement vector d. In this example we select a displacement of one pixel in the horizontal direction.

Displacement Mask

a	b

1	1	0	0
0	1	1	0
1	2	2	2
2	3	3	3

Digital Image

b \ a	0	1	2	3
0	2	3	0	0
1	3	4	1	0
2	0	1	4	1
3	0	0	1	4

Co-occurrence matrix

Figure 4.2 A simple digital image, a vector displacement mask, and the corresponding co-occurrence matrix.

Every possible combination of grey level values are entered into the mask in turn, which is then placed exhaustively at all locations in the image. Matches between the values in the mask and the grey tones in the resolutions cells of the image are counted. Once a scan is completed the total number of recorded matches is entered into the appropriate element of the matrix. the subscripts corresponding to the values within the mask.

It is usual to construct a symmetrical matrix by counting grey level pairs with a separation of both d and -d. In the example that is shown, the grey levels of 0 and 1 appear horizontally adjacent in either orientation three times. This information is represented by the value 3

in the matrix elements $\mathbf{P}_{(1,0)}$ and $\mathbf{P}_{(0,1)}$. Thus, it can be clearly seen that an element $\mathbf{P}_{(i,j)}$ is a measure of the probability of finding within the image the grey levels of i and j at the spatial separation specified by the displacement vector.

4.3.3 Practical Implementation of Matrix Construction

The scheme for matrix construction described above is perhaps the simplest to understand. However, the actual computer implementation of matrix construction did not count matches between pixel values and those of a mask as it was scanned across the image. Instead the mask was effectively 'empty' and picked up the values of the pixels at each location. These values were then used directly as subscripts to access and increment the relevant matrix element. Using this scheme only a single pass needs to be made through the image rather then one for each grey level combination. This is clearly a much more efficient implementation.

4.3.4 The Displacements Vectors Used

A total of twelve displacement vectors were used during these experiments. Consequently, twelve co-occurrence matrices were constructed for each textured image sample. The vectors were orientated in four directions; the left and right diagonals, horizontally, and vertically. Over each direction three vector magnitudes were used. These were displacements of 1, 4, and 8 pixels for the horizontal and vertical directions. For the diagonal directions the nearest equivalents of 1, $3\sqrt{2}$, and $6\sqrt{2}$ were used.

4.3.5 Matrix Normalization

Using displacement vectors of different magnitude affects the total sum of the entries in a matrix. As the magnitude increases there are fewer possible locations for the mask to occupy within an image of a given size. In order to account for this it is necessary to normalize the matrices if more than one displacement magnitude is to be used.

To achieve this the total number of possible locations for the displacement mask within the image area must be evaluated. Haralick et al. (1973) explain the simple rules needed to establish this value. Each matrix element must then be divided by this normalizing constant.

4.3.6 Reducing the Tonal Resolution of an Image

It has been stated previously that the dimensions of a co-occurrence matrix reflect the tonal resolution of the digital image. The original tonal resolution of 256 grey levels, captured with the video camera and frame store, would therefore produce very large matrices. In fact each matrix would be equivalent to a complete image in terms of the data generated. If only a small proportion of the full image is to be analysed we would expect most of the matrix elements to remain empty.

It was therefore necessary to reduce the tonal resolution in the images prior to constructing the co-occurrence matrices. This ensured that matrices of a manageable size were produced. It also eliminated wasted space and prevented excessively slow calculation times.

There is unlikely to be any advantage in using images with a tonal resolution greater than 64 grey levels unless the textured region to be analysed is exceedingly large. During these experiments the area selected for texture analysis was never larger than a very modest 64x64 pixels. A reduction of the tonal resolution down to either sixteen or thirty-two tones was therefore considered a reasonable procedure. This was performed while constructing the montage image files so that it need not be repeated for each run of the matrix production program.

4.3.7 Co-occurrence Texture Measures

Information is extracted from a co-occurrence matrix using second-order statistics to produce a numerical evaluation of the texture properties. Haralick et al. (1973) proposed fourteen texture features that could be derived from a matrix. These were all essentially measures of the concentration and the distribution of counts within the matrix. They noted that many of the proposed measures were likely to be highly correlated and thus carry similar information.

Six statistical measures, or texture features, were used during this study. These are given the symbolic names ASM, ENT, CON, VAR, IDM, and COR. The definition of these is as follows:

Angular Second Moment:

$$ASM = \sum_{i}\sum_{j} \mathbf{P}_{(i,j)}^{2}$$

Entropy:

$$ENT = -\sum_{i}\sum_{j} \mathbf{P}_{(i,j)} \log(\mathbf{P}_{(i,j)} + c)$$

Contrast:

$$CON = \sum_{i}\sum_{j} (i-j)^2 \; P_{(i,j)}$$

Variance (Sum of Squares):

$$VAR = \sum_{i}\sum_{j} (i-u)^2 \; P_{(i,j)}$$

Inverse Difference Moment:

$$IDM = \sum_{i}\sum_{j} 1/(\; 1+(i-j)^2 \;) \; P_{(i,j)}$$

Information Measure of Correlation:

$$COR = (\; 1 - \exp[\; -2.0(XY2 - XY1) \;] \;)^{1/2}$$

where,

$$XY1 = -\sum_{i}\sum_{j} P_{(i,j)} \log(\; P_{(i,j)} + c \;)$$

$$XY2 = -\sum_{i}\sum_{j} px(i)py(j) \log(\; px(i)py(j) \;)$$

px(i) = ith entry in the vector **px**, the sum of the rows of $P_{(i,j)}$

py(i) = ith entry in the vector **py**, the sum of the columns of $P_{(i,j)}$

u = the mean value of the matrix elements

c = an arbitrarily small positive constant

4.3.8 Texture Information in the Co-occurrence Matrix

Co-occurrence matrices capture information on the textural properties, or the spatial distribution of grey levels, in the image analysed. More specifically, they capture information on the grey tone transitions that exist between pixels separated by the displacement vector. To see how matrix values are affected by textural properties we can consider some of the basic perceptual properties of texture, as identified by Tamura et al. (1978).

Coarseness is a property fundamental to the perception of texture. In a coarse texture low spatial frequencies are dominant, producing large regions of low contrast. In this situation we expect mainly small transitions between the grey levels except at the boundaries between the texture primitives. This will lead to a concentration of counts around the principal diagonal of the co-occurrence matrix. As coarseness decreases more transitions between texture primitives will occur and consequently the distribution in the matrix should become more dispersed. Larger displacement vectors will also tend to select pixels that lie across primitive boundaries more frequently and so the concentration effect should be reduced in the corresponding matrices.

Contrast is another fundamental property of texture, but one that can be subdivided into 'global' contrast and 'local' contrast. For the images used in these experiments the global contrast levels are likely to be affected by external influences. For instance, the contrast setting on the SEM during the initial acquisition of the images, and the lighting level while digitizing the prints, would both affect global contrast characteristics. Variations at this global level are undesirable as they do not relate to textural properties. The histogram equalization method, described previously, was used to eliminate variations in the global contrast level.

Local contrast is a factor that does relate to textural properties. At this level we are concerned with the grey tone differences between closely spaced regions. However, the inseparable interrelationship of coarseness and contrast noted by Haralick (1979) becomes an important factor. Although an image may contain a fine texture with rapidly varying tones the level of local contrast can affect our perception of the two properties. If the local contrast within such a region is very low

the average intensity of the region can become the dominant perceived property and it will appear as a coarse texture primitive. As the local contrast is increased the rapid variations 'stand out' and we begin to perceive a very fine texture.

Relating these properties to the co-occurrence matrix we find again that the concentration around the principal diagonal an important measure. As local image contrast increases the co-occurrence counts will shift away from the principal diagonal of the matrix. Clearly, a measure of concentration along the diagonal will reflect both the coarseness and the local contrast of the texture. This description is provided by the CON measure.

The ASM measure is regarded by Haralick et al. (1973) as an indication of image homogeneity. In a homogeneous image a few grey tone transitions are dominant leading to a concentration of counts into a few matrix elements. As homogeneity decreases more varied transitions occur and the counts become dispersed to give a large number of small values in the matrix. The ASM feature, calculated as the sum of squares of the entries, will therefore take on larger values as homogeneity increases. ASM will of course be affected to some extent by the coarseness as this too leads to a concentration of values.

The correlation measure COR, corresponds to linear dependency among the grey tones. It will take on larger values in an image where grey tones are linearly related. This would be roughly equivalent to detecting long run-lengths of similar grey tones.

Other features defined by Haralick et al. relate to less tangible properties of the texture. This, in fact, is a disadvantage of most statistical approaches to texture quantification. Although Haralick's measures may provide valuable descriptions of matrix distribution it can be difficult to relate these to common perceptions of textural properties.

Another important perceptual property of texture identified by Tamura et al. (1978) is that of directionality. Haralick et al. suggested using the absolute range of the feature values over four displacement directions as an indication of directionality. In a highly directional texture we would expect the range to be large. However, the textures of the pollen taxa used in these experiments were fairly isotropic; only Pine displayed any significant degree of directionality. Therefore, the range over the vector directions was not employed for texture description during this study.

In order to account for any small degree of directionality that existed in the samples analysed, the mean value of each feature over the four vector directions was taken. In this way features that are reasonably invariant to the rotation or the orientation of the texture are produced. The mean values were calculated for each displacement size to produce a total of eighteen texture measures from each image sample.

4.4 LAWS' MASKS

4.4.1 Introduction

Texture features based on Laws' masks were introduced by Pietikainen et al. (1983). The Laws' masks define grey level patterns in local pixel neighbourhoods. The features produced are based upon the average level of correspondence between local neighbourhoods in the textured image and these predefined patterns. The construction of the masks and the extraction of texture measures is explained below.

4.4.2 Constructing Laws' Masks

The Laws' masks are all derived from three simple vectors of length 3. These are illustrated in Figure 4.3(a) and are given the identifications l3, e3 and s4. They represent the one-dimensional operations of centre-weighted averaging, edge detection, and spot detection respectively.

If these initial three vectors are convolved with each other a further five vectors of length five may be derived. The five new vectors, illustrated in Figure 4.3(b) are given the identifications l5, e5, s5, r5, and w5. The vectors l5, e5, and s5 again represent local averaging, edge and spot detection, while the r5 and w5 vectors may be regarded as ripple and wave detection schemes.

l3	(1,2,1)	e3	(-1,0,1)	s3	(-1,2,-1)

(a) Definition of Laws' three simple vectors

l5	(1, 4, 6, 4, 1)	=	l3 * l3
e5	(-1,-2, 0, 2, 1)	=	l3 * e3
s5	(-1, 0, 2, 0,-1)	=	l3 * s3
r5	(1,-4, 6,-4, 1)	=	s3 * s3
w5	(-1, 2, 0,-2, 1)	=	- e3 * s3

(b) Definition of the extended Laws vectors

Figure 4.3 Laws' vectors used to produce the texture masks.

4.4.3 Masks Used for Texture Analysis

The Laws' masks are produced by simply taking the inner product of a column and row vector, where the column vector is formed by taking the transpose. The five 3x3 masks used in these studies are defined in Figure 4.4, and the eight 5x5 masks are shown in Figure 4.5. It can be seen that all are zero-sum masks (i.e. their coefficients sum to zero).

Pietikainen et al. found that the power of the masks depended on their general form rather than the specific coefficient values used. They tried several variations of a genuine Laws mask and found that these performed about as well, or in some cases better, than the original masks. To experiment a little with this idea the two masks (d) and (e) in Figure 4.5 are modified forms based on the original Laws masks indicated.

The masks (a) and (h) represent the two-dimensional operations of center-weighted vertical and horizontal edge detection. The masks (b) and (f) may be considered as 'wave' or 'streak' detectors. A center-weighted spot detector describes the mask (d), while (g) represents

(a) E3′ * E3

1	0	-1
0	0	0
-1	0	1

(b) S3′ * S3

1	-2	1
-2	4	-2
1	-2	1

(c) L3′ * S3

-1	2	-1
-2	4	-2
-1	2	-1

(d) S3′ * E3

1	0	-1
-2	0	2
1	0	-1

(e) E3′ * S3

1	-2	1
0	0	0
-1	2	-1

Figure 4.4 The five 3x3 Laws' masks

(a) L5′ * E5

-1	-2	0	2	1
-4	-8	0	8	4
-6	-12	0	12	6
-4	-8	0	8	4
-1	-2	0	2	1

(b) L5′ * S5

-1	0	2	0	-1
-4	0	8	0	-4
-6	0	12	0	-6
-4	0	8	0	-4
-1	0	2	0	-1

(c) E5′ * S5

-1	0	2	0	-1
-2	0	4	0	-2
0	0	0	0	0
2	0	-4	0	2
1	0	-2	0	1

(d) R5′ * R5m

-2	-2	-2	-2	-2
-2	3	3	3	-2
-2	3	8	3	-2
-2	3	3	3	-2
-2	-2	-2	-2	-2

(e) S5′ * E5m

-1	-2	0	2	1
0	0	0	0	0
2	4	0	-4	-2
0	0	0	0	0
-1	-2	0	2	1

(f) S5′ * L5

-1	-4	-6	-4	-1
0	0	0	0	0
2	8	12	8	2
0	0	0	0	0
-1	-4	-6	-4	-1

(g) R5′ * R5

1	-4	6	-4	1
-4	16	-24	16	-1
6	-24	36	-24	6
-4	16	-24	16	-1
1	-4	6	-4	1

(h) E5′ * L5

-1	-4	-6	-4	-1
-2	-8	-12	-8	-2
0	0	0	0	0
2	8	12	8	2
1	4	6	4	1

Figure 4.5 The eight 5x5 Laws' masks

the two-dimensional 'ripple' detector. Finally the masks (c) and (e) are rather more difficult to name, perhaps 'chevron detectors' would be the appropriate term.

4.4.4 Texture Measures

Pietikainen et al. record that the best statistics for texture discrimination are the sums of the squared or the absolute values of the image pixels after convolution with the masks has been performed. Use of the absolute values is preferable since it reduces the amount of computation involved. This texture analysis scheme produced a total of five texture descriptors from the 3x3 masks, and eight descriptors from the 5x5 masks.

The texture features produced from this method tended to have very large values. In Chapter Six these texture measures are combined with those from the other texture analyzers which produce feature values that are several orders of magnitude smaller. This could be a potential problem since the classification stage employs a matrix inversion routine. Inverting a matrix whose elements have widely differing values can lead to excessive computational errors (Stewart (1973), Steinberg (1974)). Therefore, all the Laws' mask texture measures were divided by a suitable fixed denominator (the value of 10000 was used) so that the feature values were of comparable magnitude to the outputs from the other texture analyzers.

4.5 TEXTURE MEASURES BASED ON EDGE PAIRS

4.5.1 Introduction

The texture measures to be described in this section are based on paired edges and are a form of structural analysis rather than the statistical methods described above. In a structural approach the texture is considered to be composed of texture primitives. Texture primitives may be defined as connected regions satisfying certain properties (Wang et al., 1981), uniform grey tone being the most common property utilized. In order to quantify the texture the characteristics of the primitives must be captured, along with information on the structural or spatial organization of the primitives. This is sometimes called a syntactic approach which uses a type of 'texture language'. The syntax describes the way the texture is constructed from primitives, and the primitives may be thought of as the 'words' of the texture language.

This approach to texture analysis is initially very appealing as it relates easily to human perceptions. Despite this advantage there have been very few successful applications of structural analysis techniques. A common problem is in making the technique general enough to be used on a wide variety of image classes. Another problem is the heavy computation frequently involved in extracting texture primitives.

4.5.2 Schemes for Extracting Texture Primitives

Wang et al. (1981) describe one of the simpler schemes for primitive extraction which is based on fixed percentile thresholding. Hong et al. (1980) used an edge based extraction process that evaluated a confidence

rating for individual pixels belonging to a texture primitive. The confidence level was determined by the number of times the pixel was found to occupy a position between paired antiparallel edges. This property was considered to indicate that the pixel was an interior point of a primitive. When this computation had been completed the confidence ratings were modified by taking into consideration the ratings of the eight nearest neigbouring pixels. This was to allow for fragmented edge boundaries occurring around the primitive. Finally the values were thresholded and those pixels exceeding the threshold level were considered to form the primitives.

An interesting feature of this work was that the search for antiparallel pairs was performed on the dark side of the edges. In most textures this would be considered equivalent to measuring the properties of the areas between primitives. Although there is no reason why the lighter regions should be considered as primitives rather than the darker regions this is usually the case.

Once primitives had been isolated six properties were measured; area, perimeter, dispersiveness, elongatedness, eccentricity, direction of major axis, and average grey level. The mean and standard deviation of these values taken over all primitives larger than a threshold area were the final texture measures.

4.5.3 Measuring Texture From Paired Edges

An alternative texture analysis scheme, also based on the identification of edge pairs, was proposed by Pietikainen and Rosenfeld (1982). This has considerable similarities with the method described above, but it

avoids the need for the explicit extraction of primitives. This is an advantage as it reduces the quantity of computation involved. More importantly, Pietikainen and Rosenfeld note that textures may be perceived even if primitives are not fully defined. For example, local clusters of edges may produce a texture even if the edges do not completely surround connected regions to form primitives. This might be important for some of the pollen textures to be analysed. In particular the texture of Pine, while undeniably distinctive, is unlikely to be fully captured in primitives requiring complete boundary definitions.

The texture measures to be used are derived directly from pairs of facing edges. Again the assumption is made that these will represent opposite sides of a primitive. There is a close relationship between this method and the generalized co-occurrence matrices discussed by Davies et al. (1979, 1981). Generalized co-occurrence matrices analyse local properties that satisfy certain spatial constraint predicates. Edges are frequently used as the local property and the spatial constraint is that they occur in certain relative positions and orientations. These edge pair measures are also defined by recording edges with specific orientation and separation. They are simpler however, requiring the calculation of first-order rather than second-order statistics.

4.5.4 Preparation of the Textured Images

A number of preparation processes must be applied on the textured images before edge pairs can be identified and the texture measures calculated. These are described below.

The first stage necessary was to process each image with a low pass

filter. This was used to suppress the weaker edges, especially those that are caused by noise. Any low pass filtering operation could be used. For example, a simple mean filter or even low pass filtering in the Fourier frequency domain would be sufficient. However, the nonlinear median filter is perhaps the most appropriate, as it reduces the noise component without causing too much degradation to the stronger edges present. The images were filtered twice with a 3x3 median operator to achieve the desired effect.

The next stage was the production of an edge map for the texture samples. The eight simple 3x3 gradient filters shown in Figure 4.6 were used for detecting edges. At every location all eight filters were convolved with the image and the filter producing the strongest output selected. An alternative method would be to use a differential edge filter such as the Roberts or Sobel operators, and quantize the vector gradient direction into eight 45^{o} divisions. When completed, this operation yields an edge map indicating both the orientation and the magnitude of edges at each pixel location.

Processing was then switched to the edge map. The map was thresholded and all edge responses weaker than the threshold value discarded. The selection of the threshold level is somewhat arbitrary and should be done on a 'trial and error' basis for any given application. For this study a threshold level of 11 was set after some experimentation. This level seemed to leave edges that corresponded well to discernible features in the original textured images.

Finally, a non-maxima suppression algorithm was used to thin down the edge responses that remained. The algorithm used a 1x3 neighbourhood centered on each pixel and oriented in the direction of the edge

1	1	1
0	0	0
-1	-1	-1

0	1	1
-1	0	1
-1	-1	0

-1	0	1
-1	0	1
-1	0	1

-1	-1	0
-1	0	1
0	1	1

1	1	0
1	0	-1
0	-1	-1

1	0	-1
1	0	-1
1	0	-1

0	-1	-1
1	0	-1
1	1	0

-1	-1	-1
0	0	0
1	1	1

Figure 4.6 The eight edge templates used for edge map production

gradient at the centre pixel. If either of the neigbouring pixels had a greater edge response the centre pixel's value was suppressed by setting it to zero.

4.5.5 Detecting Edge Pairs

Edge pairs are detected by first locating each edge maxima point within the edge map in turn. A search for an antiparallel edge is initiated at each maxima point found. The search is made on through pixels that are orientated along the direction of the edge gradient at the initiating point. An antiparallel or matching edge is considered to be one that has an approximately opposite gradient direction. More specifically, if the starting pixel has an orientation of $45(i)^{o}$, we consider an edge to match if it has an orientation of $45(i+3)^{o}$, $45(i+4)^{o}$, or $45(i+5)^{o}$, modulo 360^{o}.

The search was conducted on both the light and dark sides of the initiating pixel but only to a maximum search distance. If a match had not been located within this distance the search was abandoned. The threshold distance should be set after considering the size of the

primitives that occur within the textures to be analysed. In this study it was set at 13 pixels as the dimensions of most discernible primitives were within this limit.

4.5.6 Edge Based Texture Measures

Once an edge pair had been located the primitive statistics given below were recorded. These measures were recorded separately for matching pairs found from searching on both the light and the dark side of the initiating pixel.

d The distance between the paired edges

u The mean grey level along the line connecting the edges

D The absolute difference between the contrast, or gradient magnitudes, at the two edges

We would expect measurement d to reflect the size and spacing characteristics of the primitives. Measurement u should record the tonal properties of both the primitives and the inter-primitive regions. Finally, measurement D should record how well and how consistently the primitives' boundaries are defined.

The texture features that were derived from these measures and used for classification purposes were the mean and standard deviation of d and u, and the mean value of D. In addition, the total number of paired edges, N, was employed as an extra texture descriptor. The measurements recorded on the dark and the light sides of edges were still kept separate while calculating these descriptors. A total of twelve

texture measures were therefore available from the edge based texture analyzer, six based on dark-side searches and the remainder on light-side searches.

CHAPTER FIVE

CLASSIFICATION PROCEDURES

5.1 SOME PRINCIPLES OF MULTIVARIATE CLASSIFICATION

The methods for texture analysis described in Chapter Four simply supply a numerical evaluation of the textural properties in an image sample of the exine surface. Classification rules must be developed in order to assign a sample to the correct pollen class. These rules are constructed using the texture measurements obtained on images from known pollen classes. Once they have been developed, the rules may be applied to the problem of classifying unknown image samples.

It is unlikely that any single texture measurement will be able to adequately distinguish between the pollen classes. Combinations of measurements can provide greater discriminating power, provided that they are not highly correlated. This necessitates the use of multivariate classification schemes. In these schemes we consider each feature, or texture measure in this case, to be represented by one dimension of a multidimensional or multivariate space, often referred to as the feature space.

A collection of texture measurements from samples of known class are used to form mathematical decision boundaries that divide the feature space into a number of regions. Each region is related to a specific class. In order to classify an unknown sample its texture measurements are plotted in the feature space and it is allocated to the class in whose region these measurements fall.

More formally, the feature space is partitioned into C_i regions, $i=1,2, \ldots n$, where n is the number of classes. An object k is assigned to class i if its feature vector x_k (the collection of texture measurements) lies within the region C_i (Gordon, 1981). The

boundaries between the class regions are known as the decision surfaces. There are a number of ways to construct the decision surfaces and these are considered next.

5.2 THE PARALLELEPIPED CLASSIFIER

The parallelepiped or box classifier is probably the simplest classifier design. It uses the least computer resources and is popular since it is both fast and efficient. Parallelepiped regions are constructed by specifying acceptable limits on the statistical distribution of each variable. This limit might be set, for example, as a distance of two standard deviations on either side of the arithmetic mean. Any feature vector that falls within the limits imposed upon all variables is assigned to the appropriate class. The concept is shown in Figure 5.1 for a two variable problem.

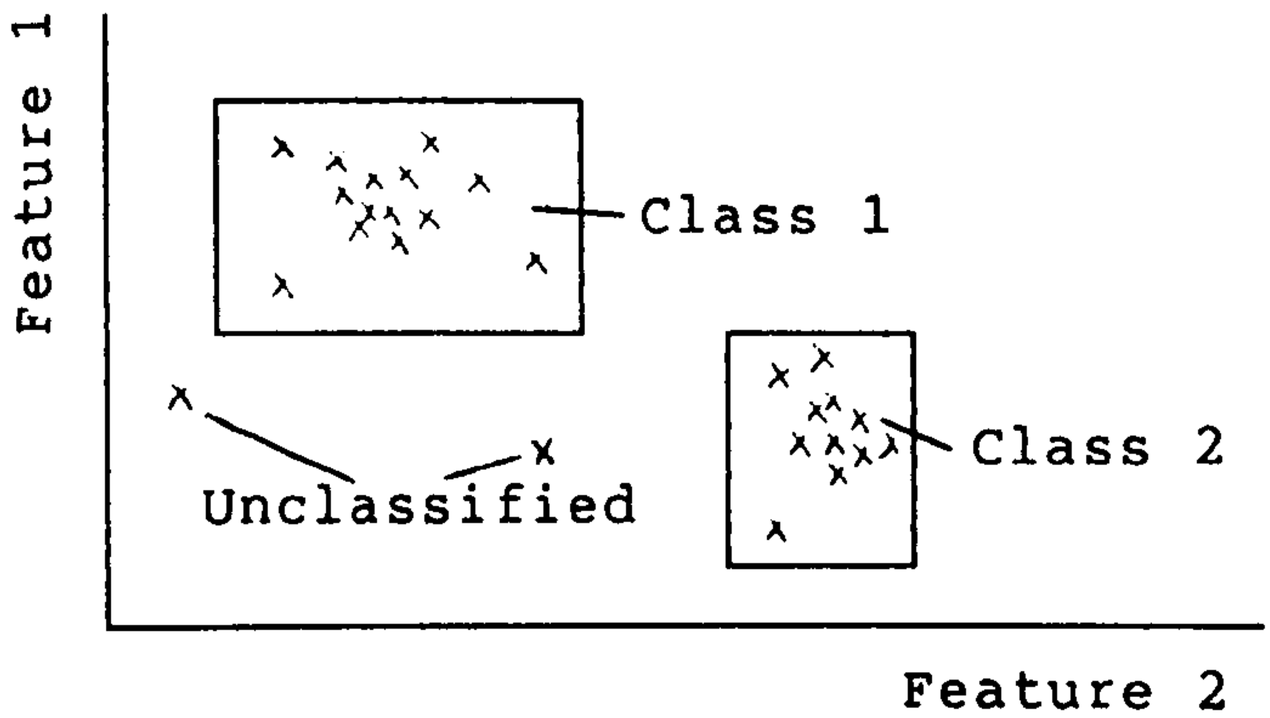

Figure 5.1 The Parallelepiped Classifier.

The major drawback with the parallelepiped classifier is that it does not calculate any form of 'certainty weighting' when classifying a sample (Thomas et al., 1987). When the feature vectors of unknown samples

are plotted in the feature space they either fall within the boundaries of a parallelepiped region, in which case they are assigned to the appropriate class, or they don't and are classified as unknown. It is as simple as that.

5.3 THE NEAREST EUCLIDEAN DISTANCE CLASSIFIER

This is still a relatively simple classification scheme which evaluates the distance, in multidimensional space, between the position of the sample vector and the location of the mean centroid vector of each class under consideration. The unknown sample is allocated to whichever class is the nearest on the basis of this measurement. The square of this distance can be expressed mathematically as,

$$d_{2i}^{2} = (\mathbf{x}_k - \mathbf{m}_i)^T * (\mathbf{x}_k - \mathbf{m}_i)$$

where $\mathbf{x}_k$ is the feature vector of sample k, $\mathbf{m}_i$ is the mean centroid vector for class i, and T represents the transpose of the vector. This is equivalent to,

$$d_{2i}^{2} = (x_1 - m_{1i})^2 + (x_2 - m_{2i})^2 + \ldots + (x_n - m_{ni})^2$$

where x is the value of each feature and m_i is the mean value of this feature for samples taken from class i.

This measure is therefore the square of the Euclidean distance between the vectors $\mathbf{x}$ and $\mathbf{m}_i$. A second measure, the so called 'city block' distance, is sometimes used for computational efficiency. This produces piecewise linear boundaries that approximate the linear decision surfaces

defined by d_2 above (Schowengerdt, 1983). The definition of the city block measure is,

$$d_{1i} = | x_1 - m_{1i} | + | x_2 - m_{2i} | + \ldots + | x_n - m_{ni} |$$

It can be shown that the Euclidean distance classifier, using the d_2 distance measure, is a special case of the maximum-likelihood classifier (Schowengerdt, 1983). It assumes that the variance-covariance matrices of each class are equal and that all a priori probabilities are equal. Furthermore, it assumes that entries in the variance-covariance matrices are confined to diagonal values. This means that all the features are uncorrelated. Finally, it assumes that each feature has equal variance.

The problem with this classification scheme is that it ignores the spread of data points in the multidimensional feature space. It operates only on independent single-dimensional representations of the probability distributions, failing to take account of the multivariate distribution and the correlations that are likely to exist between the variables.

5.4 THE MAHALANOBIS DISTANCE MEASURE

The shortcomings of the Euclidean distance classification scheme noted above may be largely overcome by using the Mahalanobis distance measure. This also evaluates the distance between the unknown sample's feature vector and the position of the mean centroid vector of each class. However, these distances are normalized with respect to the multidimensional variance of each measured variable. Thus the Mahalanobis distance, usually represented as M^2, is equal to the square of the distance expressed in units of variance for that class (Thomas

et al., 1987). It is defined as,

$$M_i^2 = (x_k - m_i)^T * S_i^{-1} * (x_k - m_i)$$

where S_i is the variance-covariance matrix for class i. This measure is a component in most maximum-likelihood classification schemes.

5.5 THE LINEAR DISCRIMINANT CLASSIFIER

5.5.1 The Fisher Linear Discriminant Function

The Fisher linear discriminant function is an alternative way of constructing decision surfaces in the feature space. Suppose we have a series of texture measurement vectors,

$$x_{1i}, x_{2i}, \ldots, x_{ni} \quad \text{and}$$
$$x_{1j}, x_{2j}, \ldots, x_{nj}$$

belonging to the classes i and j respectively. We attempt to construct a linear combination of the form

$$f_1 x_1 + f_2 x_2 + \ldots f_n x_n$$

which has values that are large for population i and small for population j, or vice versa.

Let the mean centroid vector and the variance-covariance matrix be designated as m_i, S_i and m_j, S_j for each class respectively. Using the vector maximization theorem, it may be shown that the value of f

maximizing the separation is given by,

$$\mathbf{f} = (\mathbf{S}_i + \mathbf{S}_j)^T * (\mathbf{m}_i - \mathbf{m}_j)$$

The vector **f** is known as the Fisher linear discriminant function. The linear combination presented above effectively projects a feature vector onto some linear direction or axis within the multidimensional space such that the separability of the two classes is maximized.

For the purposes of classification all that is required is to locate a class decision boundary on the projection axis. This is achieved by projecting the sample vectors that were used for constructing the discriminant function. Thus we calculate the values of $\mathbf{f}^*\mathbf{x}_i$ and $\mathbf{f}^*\mathbf{x}_j$ and construct the class boundary at the location providing the minimum classification error using the formula,

$$\text{Boundary} = (m_i * s_j + m_j * s_i) / (s_i + s_j)$$

where m_k and s_k are the mean and standard deviation of the projected values for class k.

5.5.2 A Pairwise Linear Discriminant Classifier

The Fisher linear discriminant function described above may be used as a decision rule that divides the feature space. Any class pair may be separated by an appropriate discriminant function and its associated decision boundary along the axis of projection.

The Fisher linear classifier constructed by this method is the optimal

linear classifier and yields the minimum error probability classification provided both classes have a multivariate normal distribution with equal variance-covariance matrices (Weszka et al., 1976).

The assumption of equal variance-covariance matrices allows the hyper-surfaces partitioning the feature space to be simplified to linear hyper-planes, but it means that the classifications obtained are not necessarily the maximum-likelihood classifications (Schowengerdt, 1983).

Since we have more than two classes, the classification process is repeated for all possible class pair combinations and a 'vote' is allocated to the winning class on each occasion. At the end of this procedure the sample is assigned to the class that receives the greatest number of votes.

When using a discriminant function that was constructed with the true sample class, we expect the unknown sample to be classified correctly and a vote allocated accordingly. The allocation of the vote for other 'redundant' discriminant functions should be essentially random, thus leaving the actual class with the highest vote (Weszka et al., 1976).

5.6 TESTING CLASSIFIER PERFORMANCE

5.6.1 Training Sets and Test Sets

In order to construct the decision surfaces that segment the feature space we need data obtained from samples of known class. In the Euclidean distance classifier these measures are used to provide the mean centroid vectors of each class. In the linear discriminant classifier they

allow the construction of the Fisher discriminant function vectors and the determination of a decision boundary along the projection axis.

The data used for constructing a classifier is known as the training set. It is important that there are sufficient data in this population sample to allow accurate estimates of the true probability density functions (abbreviated to PDFs) of each class. Of course, the larger the sample size the more accurate the PDF estimates should become. Large training sets are therefore a desirable feature when constructing a classifier.

Once a classifier has been constructed we need to test its performance. The simplest method is to use the feature vectors of each sample from the training set, and treat them as if they are the feature vectors of unknown samples. However, using the same data for both training and testing can lead to over-optimistic estimates of the classifier success rate. This is especially true if the data set used to construct the classifier is small.

The reason is simply that the training set is only a population sample. Therefore only an approximation of the true PDF of the classes is estimated. Any estimate is likely to possess irregularities compared to the true population PDFs. Such differences are not accounted for within the design of the classifier, which will simply attempt to optimize the classification success of the training set presented to it. The PDFs of a new set of samples, although drawn from the same population, are likely to differ slightly from those of the training set. Consequently, the classifier will no longer be the optimal classifier for these data, and the success rate will deteriorate.

Therefore, using the same data as both a training set and a test set will lead to a favourably biased estimate of the classifier performance. The estimate of classifier performance obtained from the training set is known as the apparent error rate, while that obtained from an independent test set is called the true error rate. To obtain a true estimate an independent test set should be used.

Dividing the available data into two sets and using one for training and the other for testing would appear to be the sensible solution. However, if the amount of data available is limited, as is frequently the case, splitting it into two groups does not make best use of it (Hand, 1981). As noted above, it is desirable to have the largest possible training set in order to accurately reflect the population characteristics, but we also want a large test set when assessing the classifier performance. The leave-one-out or jackknifed classification strategy provides a solution to this dilemma.

5.6.2 The Leave-one-out Strategy

The leave-one-out technique is a simple and elegant solution to the problem of obtaining unbiased estimates of classifier performance when only a limited data set is available. Each sample in the data set is classified in turn, using a classifier constructed from all the remaining samples. This ensures that the training set is always as large as possible, yet eventually all the samples are independently classified to give an accurate assessment of the true success rate.

The penalty to pay for this attractive scheme is the increased computation involved in constructing a new classifier for each sample.

It is usually possible, with some care, to limit the amount of recalculation necessary. For instance, in the Euclidean distance classifier, 'running mean' centroid vectors may be used greatly to reduce the work involved. Similar techniques may be found for the pairwise linear discriminant classifier. By using efficient algorithms the need to restructure the classifier for each sample should not become too great a burden.

5.7 FEATURE SELECTION

5.7.1 Selecting Suitable Features for Classification

During any identification process there will be an almost infinite number of features that could be measured in order to facilitate the classification of the objects under consideration. Obviously some assessment of the 'goodness' of features is needed so that those with higher ratings can be employed in preference to the rest.

There are a number of aspects to consider when evaluating how good a feature is for classification purposes (Pankhurst, 1978). Ease of observation, for example, may be considered a valid criterion. Almost certainly the reliability of the feature will be critical in this assessment. Finally, the diagnostic power of the feature is bound to be an essential factor.

To a large extent many of these decisions have already been made in this particular application. Texture analysis on the exine was selected, in preference to the analysis of pollen shape, principally because it was considered to be a more reliable feature within the SEM environment.

Counting apertures would have only limited diagnostic power because we cannot be sure that all apertures are visible when they are recorded on a single SEM image. Having chosen texture analysis as the principal feature extraction technique, the specific methods employed were selected with 'ease of observation' borne in mind. Obviously, in this situation, ease of observation equates to computational considerations such as the time and memory required to calculate a texture measurement.

Thus, the features that are to be supplied to the statistical classifiers have already been vetted and should be those that are likely to prove the most suitable in terms of reliability, diagnostic power, and so on. Yet the question of which of the computed features are the best in a statistical classification scheme, that is which will give the best results, still remains to be answered.

5.7.2 The Need to Select Optimal Subsets

At this point it can be tempting to wonder why we should bother to select optimal variable subsets. Although it is clear that the use of several highly correlated variables can add only a little information to that already available from just one, they are still adding information. The use of more variables cannot subtract discriminating evidence, so it might seem sensible to use all the data available (Hand, 1981).

There are a number of reasons why finding the best subset of variables is still important. The first of these is that it is simply inefficient to go to the trouble of programming in feature extraction routines, performing the necessary computations and storing the data if the variables are not needed. The use of unnecessary variables will

obviously increase the time taken to perform the classification process. Finally, and perhaps rather surprisingly, the use of a greater number of variables can lead to a deterioration in the classifier performance.

Increasing the number of variables in the classifier produces a corresponding increase in the dimensionality of the feature space. Unless the size of the training set is also enlarged there is a danger that poorly sampled estimates of the PDFs will be generated as the data becomes more dispersed within the multidimensional space. This can weaken the foundations on which the classification scheme is based. Furthermore, the complexity of the decision surfaces is also increased which can make the classifier too inflexible. Although the apparent error rate estimated with the training set may decrease, an independent test set will show that the true error rate has increased due to this inflexibility (Hand, 1981).

5.7.3 Optimal Subset Selection Procedures

Optimal subset selection is concerned with finding a subset d' of variables, from the complete set d, that gives the greatest classification success rate. The simplest idea for the selection of an optimal subset is to combine the d' features that performed best when used individually. Unfortunately this is not guaranteed to find the best subset since it ignores the multivariate relationships that exist between variables. One important relationship is the correlation among the variables. Combining d' variables that are highly correlated would be little better than using a single variable alone (Hand, 1981).

The next simplest idea is the exhaustive search method. All the

possible combinations of d' from d are tried. The best combination is guaranteed to be found, but the method is only viable if the number of variables involved are small. The number of possible subsets of r variables from a total of n variables is given by the combinations rule,

$$ {}^{n}C_{r} = \frac{n!}{(n-r)!\ r!} $$

Thus, even the fairly modest request of finding the optimal five variable subset from a total of fifteen variables would require the classification procedure to be repeated over 3000 times!

There are a series of so-called 'suboptimal' search methods that are fast relative to the exhaustive search, but are not guaranteed to find the best solution. The sequential backwards elimination method is one example. Starting with the full variable set, each variable in turn is omitted and a separability index (see Section 5.7.4 below) is calculated. When this has been completed the variable which caused the least deterioration in the separability index is discarded and the process is repeated on the subset that remains. This procedure is much faster than the exhaustive search. For example, finding the subset of five variables from fifteen would require the classification process to be repeated only 105 times.

Once a variable has been discarded in this scheme it cannot be reconsidered. The relationships between discarded variables are ignored and it is this fact which can lead to the suboptimal selection of subsets.

A similar method is the sequential forward addition algorithm. As its name suggests, this is effectively the reverse process of sequential backward elimination. Starting with the best single variable, each

variable in turn is added and the one producing the greatest improvement in the separability index is kept. This algorithm can require fewer calculations if small subsets are being searched for, but it suffers from the same problems as the backward elimination method. Since the relationships between the discarded variables, or those that are still to be selected, are ignored, the two methods can produce different results.

A type of hybrid scheme is the even more strangely named 'plus-p, take-away-q' selection algorithm (Hand, 1981). This iterates through the process of adding the best combination of p variables and then subtracting the worst q variables. Values of (p=2 , q=1) are a popular selection.

5.7.4 A Separability Index

In the section above there was talk of a separability index that was used to estimate how well a set of variables could separate or distinguish classes. Several such measures are available, as outlined by Thomas et al. (1987), but the measure to be used here is the Hotelling's T^2 statistic. This a measure of the distance between group means relative to the dispersion within the samples. It is the multivariate equivalent of the well known univariate T-test statistic. The T^2 statistic is defined by,

$$T^2 = \frac{N_1 N_2}{N} (x_1 - x_2)^T * S^{-1} * (x_1 - x_2)$$

where x_i is the mean centroid vector of class i, S is the pooled

common variance-covariance matrix, N_k is the number of samples in class k, and N is the total number of samples $(N_1 + N_2)$.

The separability measure was used in this work for subset selection using the sequential backward elimination method. However, the coherence property of this statistic also allows an accelerated search algorithm to be used (Hand 1981). In the accelerated search all possible subset combinations are considered without the need to explicitly compute each one.

5.7.5 The Accelerated Search Algorithm

The coherence property of Hotelling's index states that the T^2 value for a subset is always less than, or equal to, the T^2 value of the complete set from which the subset is derived. The way the coherence property of the T^2 measure is used for an accelerated search is best explained by the use of a simple example, and the corresponding diagram drawn in Figure 5.2.

Suppose we have a two class problem with six variables measured on samples from each class. The first stage is to calculate the T^2 statistic using the full six variables which, in this example, gives a value of 379. We now proceed to construct a treelike structure by the successive deletion of variables.

The second level of the tree consists of six nodes, each of which is formed by the omission of a single variable. At each node the new T^2 statistic is calculated. The subset that maintains the highest separability index is then selected. In the diagram it can be seen that this is the

subset containing the variables {1,2,3,5,6}, which has a T^2 value of 354.

We continue by constructing the next level down on this selected branch. This provides a further five nodes each with a corresponding T^2 value. Again the subset with the largest T^2 value is selected. This is the subset {1,2,3,6} which has a separability index of 179.

So far we have followed the procedure of the sequential backward elimination method. This method would continue to pick the best subset at each level and continue to move on down the branch of the tree thus selected. In the accelerated search we now consider what other four variable subsets are possible using the variables discarded so far.

To do this the T^2 value of our best four variable subset is compared to the T^2 values obtained on the preceeding level. Starting from the left in the diagram we can see that our value of 179 exceeds the first two values of 76 and 29. It also exceeds the 130 value obtained from the rightmost subset on this level. Therefore, following on from the coherence property, we know that there can be no four variable subset derived from these sets that will exceed our maximum value of 179 found so far.

The third subset on level one had a value of 215. It is therefore possible that this may be able to yield a higher scoring four variable subset and should be investigated. It must be noted at this point that there is the possibility of considerable repetition in the tree structure. For example, the subset {2,4,5,6} may be derived from this set now under consideration. However, the same combination has already been

'rejected' as a possibility because it may be derived from the first set on level one. In fact it can be shown that a subset must contain the variables 1, 2, 4, and 6 in order for it not to have been derivable from a set already considered and rejected. Only the {1,2,4,6} subset needs to be examined and this turns out to have a lower value than the maximum of 179 found so far. At this stage we have considered, and found the best of, all four variable subsets that are possible without the need to evaluate explicitly every one.

Proceeding to the next level we find that the best three variable subset has a T^2 value of 132. Again this exceeds the values on the first level of 76, 29, and 130. Therefore any subset would have to contain the variables 1, 2, 4, and 6, in order for it not to have already been considered or rejected. This is clearly impossible for a set of three variables so we may conclude that no better subset can exist. Finally, we arrive at the best two variable combination {1,2} that has a separability index of 79.

In this particular example a sequential backward elimination would have produced the same optimal subsets at each level. However, the advantage of the accelerated search in being able to find true optimal subsets without the computational costs of an exhaustive search are clearly demonstrated. If the {1,2,4,6} combination on level two had had a larger value than the {1,2,3,6} combination we could have jumped across to this branch of the tree and continued the search from that node. The algorithm is often said to be a 'search-and-bound' technique because of this ability to skip to other nodes whenever it is appropriate.

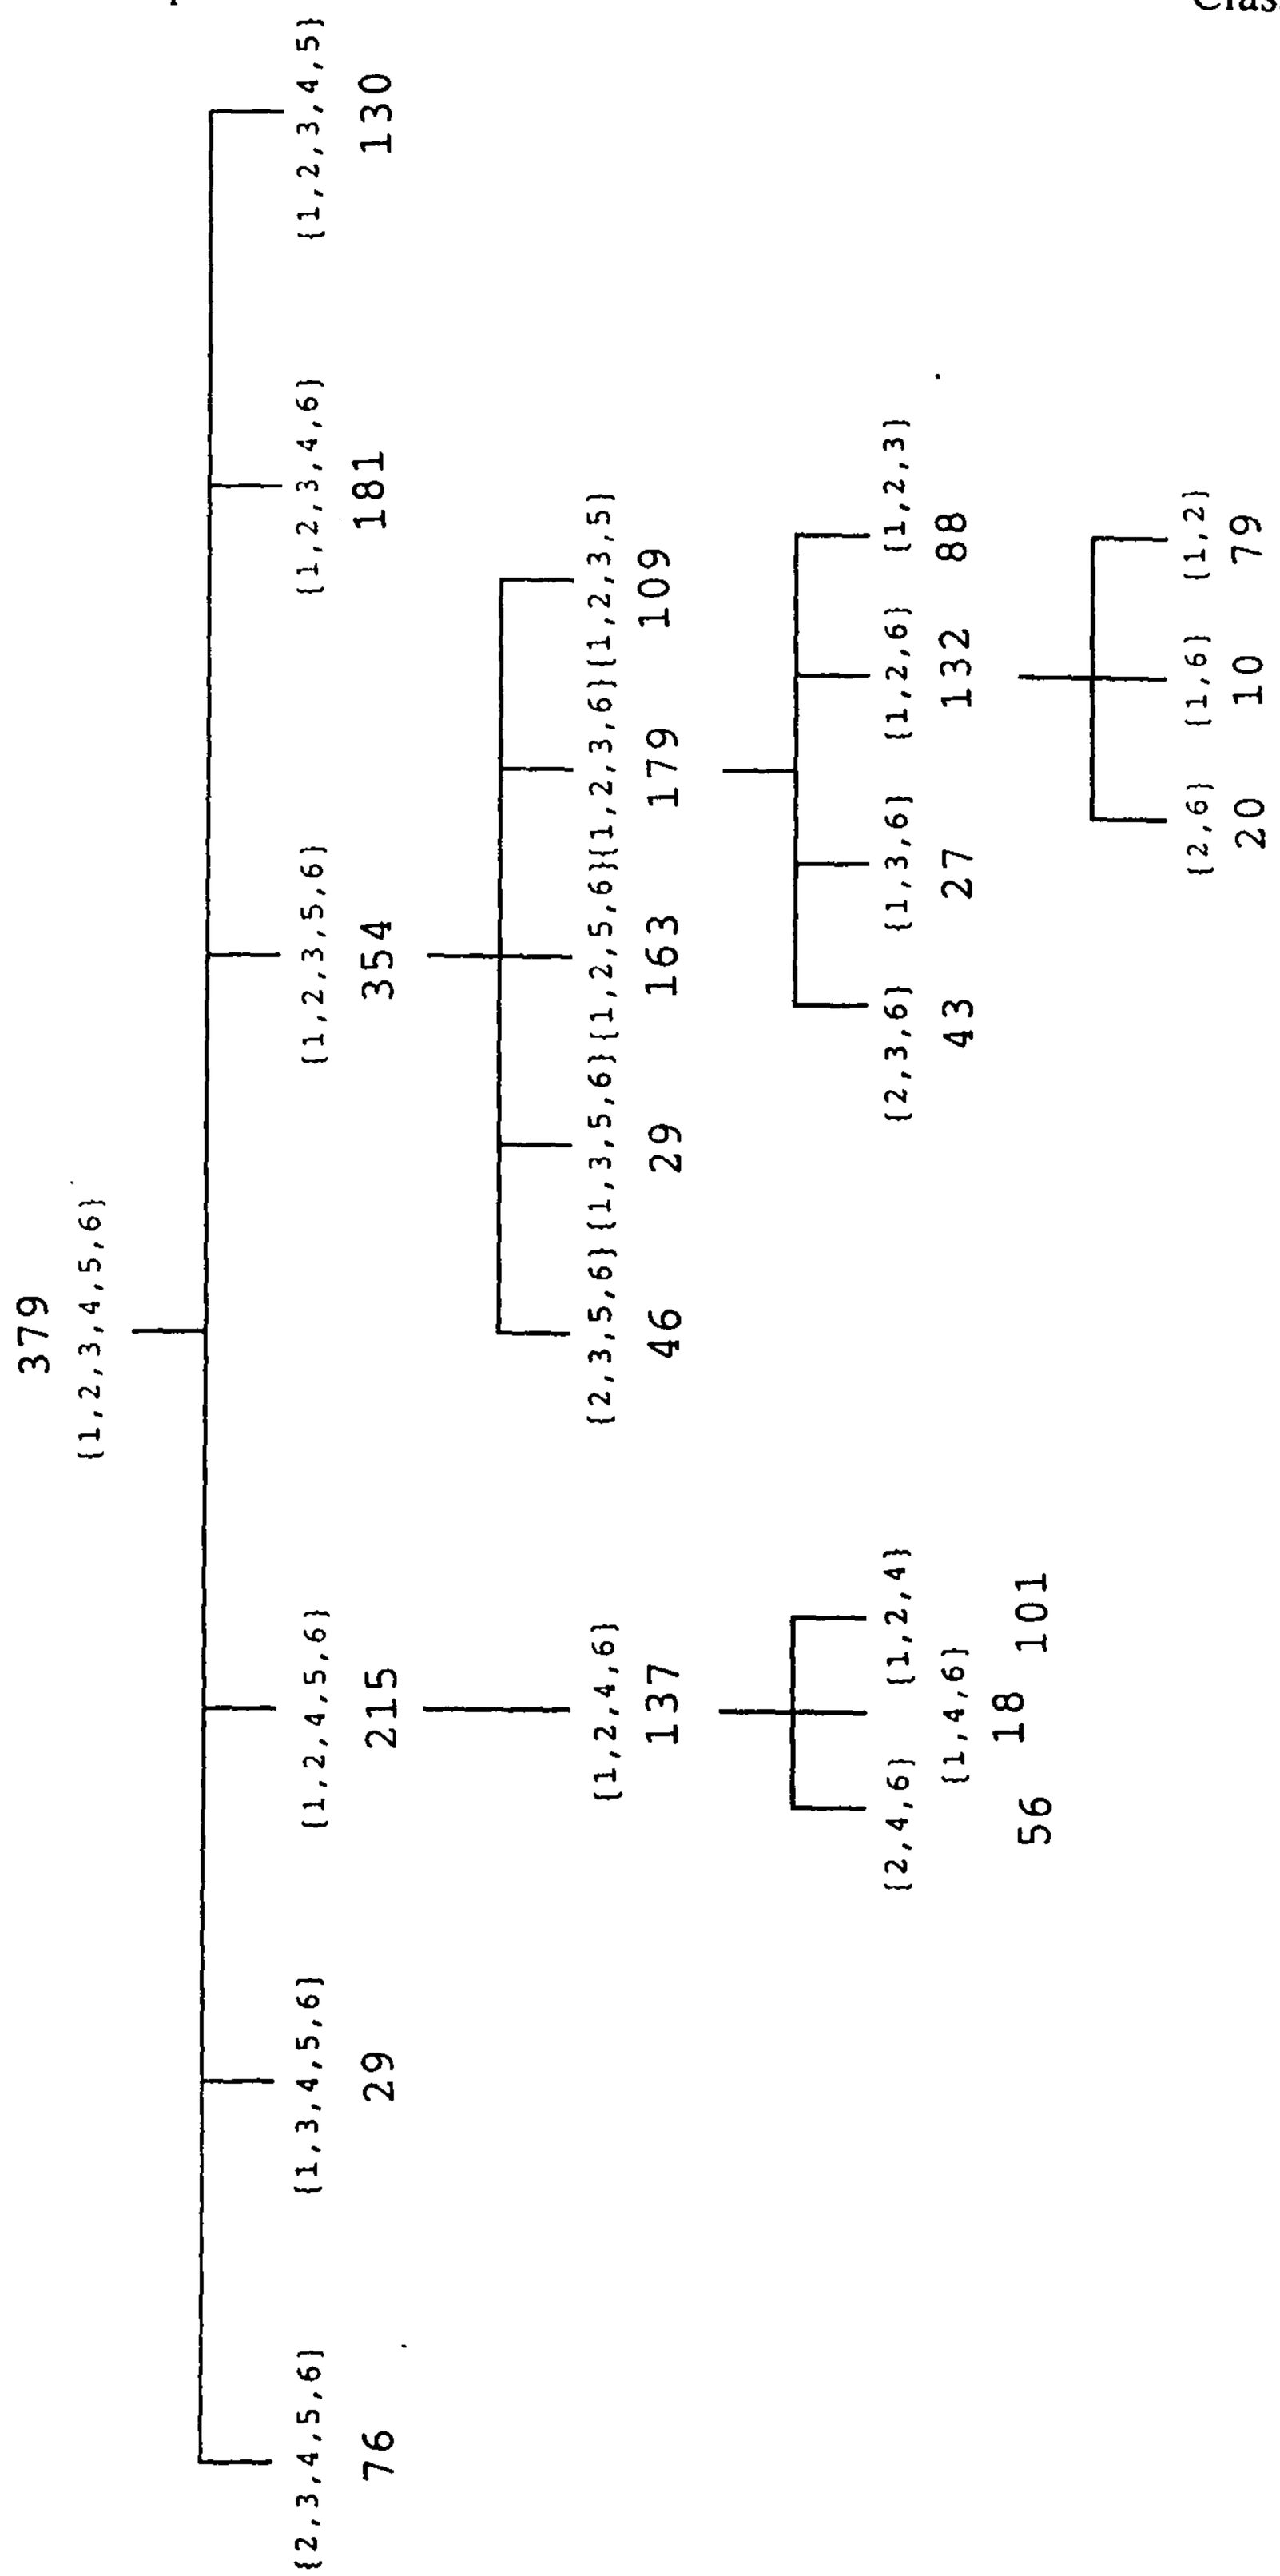

Figure 5.2 An example of an accelerated search for optimal subsets

CHAPTER SIX

CLASSIFICATION RESULTS

6.1 INITIAL EXPERIMENTS

The initial classification experiments were performed with only three pollen taxa. These were Elm, Plantain, and Hazel. Three digitized images displaying a single complete pollen grain were available for each of the taxa. Several methods of texture evaluation were investigated including those based on invariant moments (Sadjadi and Hall, 1978, Wong and Hall, 1978), and the coarseness and contrast indices described by Tamura et al. (1978). None of these appeared to be promising. In contrast, the initial results obtained from grey level co-occurrence analysis were encouraging.

Texture analysis was performed on 64x64 pixel areas selected with a 'roaming window' which could be positioned manually anywhere within the image. Once a textured patch of the exine had been selected, the tonal resolution within this region was linearly reduced to sixteen grey levels and the twelve co-occurrence matrices constructed. Only two texture descriptors, namely ASM and ENT, were extracted from each matrix. These measurements were averaged over the four displacement vector directions to produce a final total of six texture measures per sample.

In order to construct the Fisher linear discriminant functions used in the classification scheme eight texture samples were extracted for each pollen taxon and used as a training set. A further eight samples were then taken from each class, at different locations on the pollen grains, in order to form an independent test set.

When the texture measures calculated from the test set were entered into the classifier nineteen samples, or almost 80%, were correctly identified.

This was considered an encouraging start, especially since several of the test samples were positioned close to the pollen margins where the texture becomes distorted due to the curvature of the grains. Due to these promising results further investigations into co-occurrence analysis were undertaken and these are discussed next.

6.2 CO-OCCURRENCE ANALYSIS

6.2.1 Details of the Database and Texture Measures

Further experiments with co-occurrence analysis were performed on all six pollen taxa; Oak, Pine and Rye grass being added to the original three classes. Ten SEM photomicrographs of pollen grains were available for each taxon. Sub-scenes of 60x60 pixels were extracted from the digitized photomicrographs and reassembled into new image files. A complete montage file containing sixteen samples was produced for each pollen class. These files contained eight unique samples whose tonal resolution had been linearly reduced to sixteen grey levels. The remaining eight samples were duplicates except that the histogram equalization technique was employed to normalize the sub-images before reducing the tonal resolution to sixteen levels.

The number of texture measurements extracted from the co-occurrence matrices was increased in this second stage of experiments. The three measures ASM, ENT and COR are invariant under monotonic grey tone transformations (Haralick et al., 1973) and could be used with any of the sub-scenes. A second set of features consisting of the CON, VAR, and IDM measures were suitable for use only on the histogram equalized samples.

6.2.2 Classification of the Unequalized Samples

The Fisher linear discriminant classification procedure was performed first on the samples that were not histogram equalized. This restricted the number of texture features that could be used to the nine derived from the ASM, ENT and COR measures. The results of the classification are summarized in the confusion matrices presented in Tables 6.1(a-c).

A confusion matrix shows the classification distribution among the pollen classes. For samples obtained from a given taxon, represented by a row of the matrix, the figures show the proportion that are classified as any other taxon, given by the columns of the matrix. Thus, if every sample was correctly classified, all entries would be concentrated on the principal diagonal of the matrix. The codes identifying the pollen taxa are the same as those given in Section 3.1 of this report. The code Uc indicates samples that were unclassified.

Table 6.1(a) illustrates the results that were obtained from using a single texture measure. These results have been averaged over the nine texture measures that were used. The matrix indicates that considerable variation in the success rate among the pollen taxa existed. Elm was surprisingly successful (over 66% correctly classified), while Rye grass and Hazel did rather poorly (less than 20%). By analysing the raw results it was found that texture features based on vector displacements of 1 and 4 pixels averaged a 34% success rate, while those based on the 8 pixel displacement averaged only 29%. It was also noticed that the performance of the COR texture measures was significantly poorer than those of ASM and ENT.

The overall success rate using a single texture measure was 32.6%

correctly classified. Although this may appear to be rather poor it is still considerably better than that expected from chance alone (17%).

The classification performance obtained with two texture measures is illustrated in Table 6.1(b). These results are the averages taken over all thirty-six possible combinations of two features. It is immediately apparent that a large increase in the number of correct classifications has occurred. The overall average rose to 57.5% but considerable variations between pollen classes still existed. Pine and Elm did rather well while Rye grass and Hazel still lagged behind. The best pair of features managed to correctly classify 75% of the samples, while the worst pair achieved only 42%. A small number of samples (3% of the total) remained unclassified since they had an equal number of classification votes for two or more of the classes. These are represented by the Uc column of the matrix.

As more texture measures are incorporated into the classification scheme the number of feature combinations possible becomes prohibitively large. For experiments using four texture descriptors eight combinations were selected, somewhat arbitrarily, based only on intuition gained through the study of earlier results. The performance of these feature sets are summarized in Table 6.1(c).

The average classification rate rose to 89%, with a range of 94% for the best combination selected, to 79% for the worst. Again some samples, particularly those of Hazel, remained unclassified. Finally, utilizing all nine texture measures only one sample was identified incorrectly, and one remained unclassified.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl
Pi	**33.3**	8.3	16.7	26.4	0.0	15.6
El	4.2	**66.7**	0.0	1.4	11.1	16.7
Ry	9.7	6.9	**18.0**	34.7	19.4	11.1
Oa	11.1	1.4	15.3	**38.9**	15.3	18.1
Ha	5.6	15.3	31.9	9.7	**16.7**	20.8
Pl	12.5	6.9	19.4	5.6	33.3	**22.2**

Table 6.1a Confusion matrix for unequalized samples classified by a single texture measure. Figures are percentages.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**80.0**	0.7	1.7	11.8	0.7	3.1	1.7
El	1.0	**87.0**	1.0	0.3	4.2	3.1	3.1
Ry	1.7	11.8	**35.8**	16.3	19.4	7.9	6.9
Oa	13.2	0.3	5.7	**58.2**	9.6	11.0	2.5
Ha	1.7	7.9	23.5	11.1	**31.1**	21.8	2.8
Pl	5.6	0.3	13.2	9.4	17.7	**52.1**	1.7

Table 6.1b Confusion matrix for unequalized samples classified by two texture measures. Figures are percentages.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**100.0**	0.0	0.0	0.0	0.0	0.0	0.0
El	0.0	**100.0**	0.0	0.0	0.0	0.0	0.0
Ry	0.0	1.6	**82.3**	0.0	10.9	1.6	3.1
Oa	0.0	0.0	0.0	**90.6**	6.3	3.1	0.0
Ha	0.0	0.0	6.3	4.7	**76.6**	1.6	10.9
Pl	0.0	0.0	6.3	9.4	0.0	**84.4**	0.0

Table 6.1c Confusion matrix for unequalized samples classified by four texture measures. Figures are percentages.

6.2.3 Classification of the Equalized Samples

For comparative purposes the classification procedure described above was repeated for the histogram equalized samples. The corresponding confusion matrices are displayed in Tables 6.2(a-c). Studying the results for a single texture measure, shown in Table 6.2(a), it is evident that some differences exist. The performance of Pine and Rye grass are greatly improved, Hazel and Plantain are similar, but Oak and particularly Elm have deteriorated. The grand average is almost identical however at 32%. The COR measures were still found to be notably inferior to ASM and ENT.

With classification by two features, shown in Table 6.2(b), the grand average obtained (55%) was also similar to the previous result. There was, however, slightly less variation in the results between pollen taxa. More samples than previous were drawn on votes and could not be classified.

The average success rate with four texture measures, shown in Table 6.2(c), had deteriorated slightly at 80%, but using all features only a single image sample was classified incorrectly.

On the basis of these results it was felt that although histogram equalization produced quite drastic changes in the visual appearance of the samples, it did not appear to have any ill-effects on the performance of the ASM, ENT, and COR texture measures.

The advantage of using the equalized images is that a greater number of texture measurements may be extracted from the co-occurrence matrices. These additional features could not be used before since they

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl
Pi	**61.1**	0.0	22.2	2.7	6.9	6.9
El	0.0	**33.3**	1.4	11.1	22.2	31.9
Ry	16.7	5.5	**31.9**	0.0	33.3	12.5
Oa	4.2	13.8	6.9	**25.0**	18.1	31.9
Ha	5.6	12.5	26.4	12.5	**15.3**	27.8
Pl	9.7	23.6	11.1	18.1	34.7	**30.6**

Table 6.2a Confusion matrix for equalized samples classified using a single texture measure. Figures are percentages.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**87.8**	0.7	7.5	0.4	0.4	0.7	2.5
El	0.0	**52.1**	3.2	7.5	16.4	12.8	7.9
Ry	5.7	12.5	**51.2**	2.1	17.8	5.7	5.0
Oa	0.4	6.1	6.4	**62.9**	7.1	8.9	8.2
Ha	2.1	15.7	24.2	6.4	**24.5**	21.4	5.7
Pl	1.4	13.4	9.9	9.9	6.7	**52.3**	6.4

Table 6.2b Confusion matrix for equalized samples classified by two texture measures. Figures are percentages.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**100.0**	0.0	0.0	0.0	0.0	0.0	0.0
El	0.0	**78.1**	9.4	0.0	12.5	0.0	0.0
Ry	0.0	20.3	**62.5**	6.3	3.1	0.0	7.8
Oa	0.0	0.0	0.0	**100.0**	0.0	0.0	0.0
Ha	0.0	7.8	12.5	0.0	**71.8**	3.1	4.7
Pl	0.0	3.1	4.7	0.0	6.3	**82.8**	3.1

Table 6.2c Confusion matrix for equalized samples classified by four texture measures. Figures are percentages.

are sensitive to the average brightness and global contrast of the image. These are parameters that could not be accurately controlled during image acquisition on the SEM or at the digitization stage. The classification procedures were repeated on the equalized samples using the newly available features CON, VAR and IDM. The outcome is summarized in Tables 6.3(a-b).

There were no surprises with the results using a single texture measure. The overall average of 30% was in accordance with previous results, and the success of individual pollen taxa closely resembled those in Table 6.2(a). The IDM features produced the poorest results so far (averaging 22%) and features using the eight pixel displacement did badly again (averaging 21%). Hence, the IDM feature with eight pixel displacement produced a very dismal 10%.

With two features the results were not as good as those obtained before, averaging only 43%. An important discovery at this stage was that feature pairs that included one of the IDM measurements were giving the best performances (the combination of IDM at one and four pixel displacements produced a 66% success rate). This clearly illustrated that the performance of individual features could not be used to predict successful combinations.

The trend of poorer overall performance relative to the ASM, ENT and COR measures continued as larger feature sets were used.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl
Pi	**61.1**	6.9	9.7	11.1	6.9	4.2
El	9.7	**30.6**	11.1	15.3	9.7	23.6
Ry	18.0	13.8	**26.3**	12.5	18.0	11.1
Oa	2.7	12.5	30.6	**30.6**	2.8	20.8
Ha	5.6	25.0	27.7	9.7	**11.1**	20.8
Pl	8.3	27.7	25.0	9.7	6.9	**20.8**

Table 6.3a Confusion matrix for equalized samples classified using a single CON, VAR, or IDM texture measure.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**85.7**	4.2	3.1	0.7	1.0	1.7	3.5
El	4.2	**39.6**	9.7	16.7	8.0	17.4	4.5
Ry	7.6	10.1	**37.5**	14.9	12.5	6.6	10.8
Oa	0.7	9.4	17.4	**48.6**	4.9	16.0	3.1
Ha	1.0	24.0	29.2	6.9	**17.0**	19.4	2.4
Pl	3.8	21.2	17.7	16.7	6.9	**29.9**	3.8

Table 6.3b Confusion matrix for equalized samples classified by two texture measures taken from the CON, VAR, and IDM statistics.

6.2.4 Leave-one-out Classification

It was indicated in Chapter Five that classification success rates assessed by using the same data for both training and testing the classifier are vulnerable to producing over-optimistic results. In order to assess the extent of this problem the original classifier was restructured to allow leave-one-out, or jackknifed, classification. The classification of the equalized samples by the ASM, ENT, and COR features was then repeated using the leave-one-out scheme.

The results, summarized in Tables 6.4(a-b), indicate a drop in success rate, as might be expected. Employing two texture descriptors an average of 47.3% was achieved. The best pair of texture measures obtained 62% and the worst pair 31%. Using the combinations of four features an average of 58% were correctly classified, and with all features 60%.

Clearly, the previous classification scheme was producing over-optimistic estimates of the success rate. Furthermore, this degree of over-optimism increased greatly as more texture descriptors were utilised. Leave-one-out classification was therefore considered an important procedure, in the absence of independent training and test sets, allowing an unbiased assessment of the true classification error rate.

6.3 FURTHER CO-OCCURRENCE ANALYSIS

6.3.1 Details of the Database

The leave-one-out classifier avoids the problem of over optimistic classification success rates. However, a large training set is still an advantage as it should reflect the population characteristics more accurately. The classifier is then less likely to become specialised towards peculiarities present in the training set and hence it may perform better with new independent test sets.

It is also desirable when attempting to maximize classification success occasionally to renew the samples. This prevents alterations being made which, although producing better classification results, are in fact merely 'fine-tuning' the classifier to the peculiarities of the samples available.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**84.4**	0.7	9.4	1.7	0.3	1.7	1.7
El	0.0	**47.9**	4.2	9.0	17.0	14.6	7.3
Ry	6.9	13.2	**41.7**	3.1	21.2	5.6	8.3
Oa	0.7	7.3	6.9	**51.7**	6.9	12.8	13.5
Ha	2.1	16.3	27.8	6.9	**17.7**	21.5	7.6
Pl	1.4	17.7	11.5	10.4	8.3	**41.6**	9.0

Table 6.4a Equalized samples classified by a leave-one-out scheme with two texture measures drawn from ASM, ENT, and COR.

Actual Class	Classified As Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**94.6**	1.8	1.8	0.0	0.0	0.0	1.8
El	0.0	**58.9**	19.6	1.8	8.9	7.1	3.6
Ry	0.0	25.0	**33.9**	7.1	12.5	7.1	14.3
Oa	0.0	1.8	5.4	**76.8**	3.6	7.1	5.4
Ha	0.0	21.4	26.8	0.0	**28.6**	12.5	10.7
Pl	0.0	3.6	8.9	5.4	14.3	**55.4**	12.5

Table 6.4b Equalized samples classified by a leave-one-out scheme with four texture measures drawn from ASM, ENT, and COR.

Poorer performance on a new set of samples drawn from the same population may be the outcome of such actions (Hand, 1981).

It was for these reasons that a new and larger set of exine texture sub-scenes were obtained. For each of the six pollen taxa two montage image files were constructed in a similar manner to before. All samples were 64x64 pixels in dimension and were histogram equalized before reducing the tonal resolution to sixteen grey levels. Thus, the new database contained a total of 192 new exine samples, 32 unique

samples from each pollen class. The first sixteen samples of each taxon are shown in Plates 6-I to 6-VI.

6.3.2 Texture Measures

Due to the poor performance of the COR features and their relatively expensive computation time it was decided to no longer extract these from the co-occurrence matrices. The remaining five features were unaltered except that the two previously separate groups were combined to produce a single output listing. Hence fifteen texture measures were calculated and available for use from each exine texture sample.

6.3.3 Initial Classification Results

Following the procedure employed previously the samples were classified by the leave-one-out scheme using all possible pairs (now 105) of texture measures. The average rate of correct classification was 25%, and the best pair classified only 36% correctly. These results were rather poor and very disappointing considering those achieved with the previous data set. Closer inspection of the results revealed the reason for this deterioration. Many samples were remaining unclassified since they were drawn on votes between two or more classes. In fact the majority of these unclassified samples were drawn on votes between only two classes.

To alleviate this problem the classification algorithm was redesigned again. In this newly modified algorithm, whenever a sample was found to be drawn on votes between two classes its feature vector was

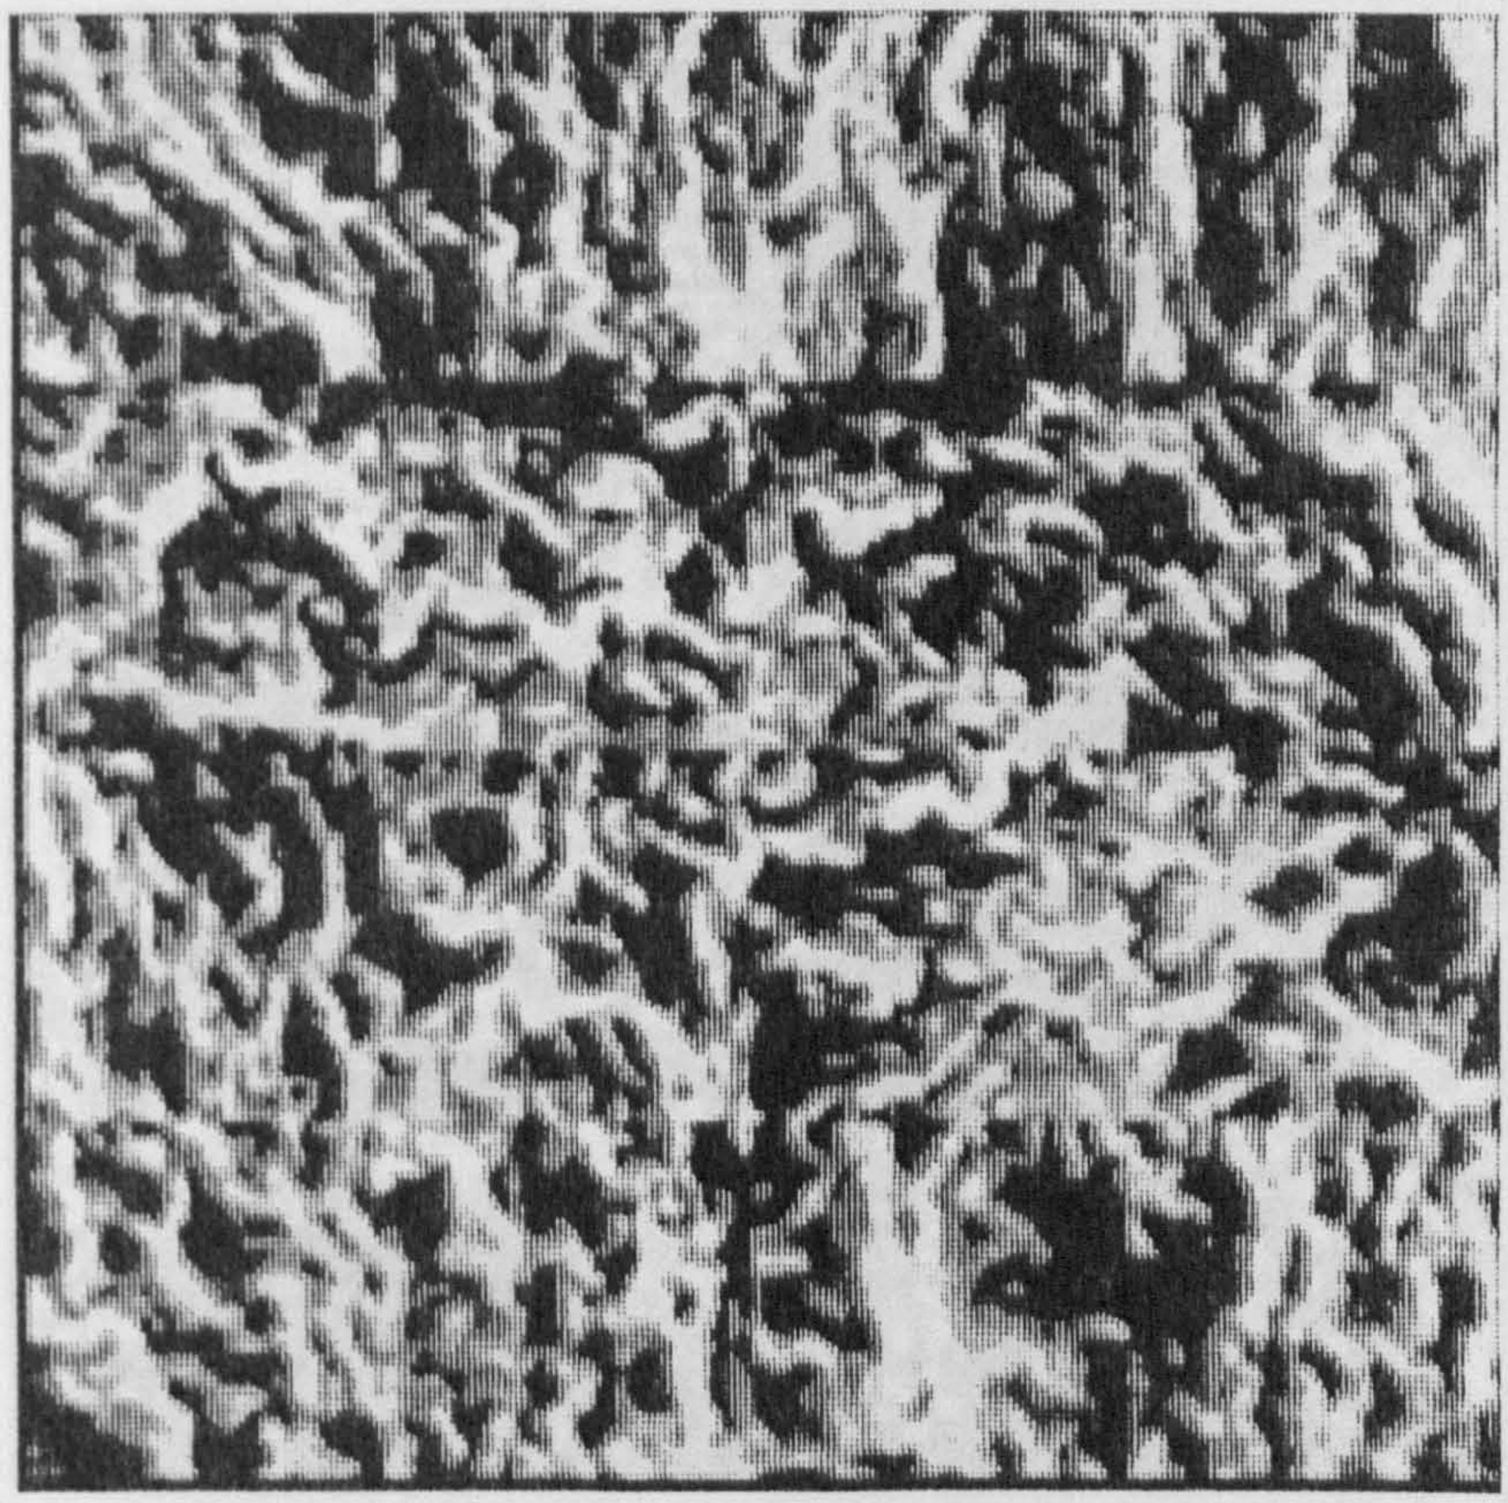

Plate 6-I. Pine exine samples.

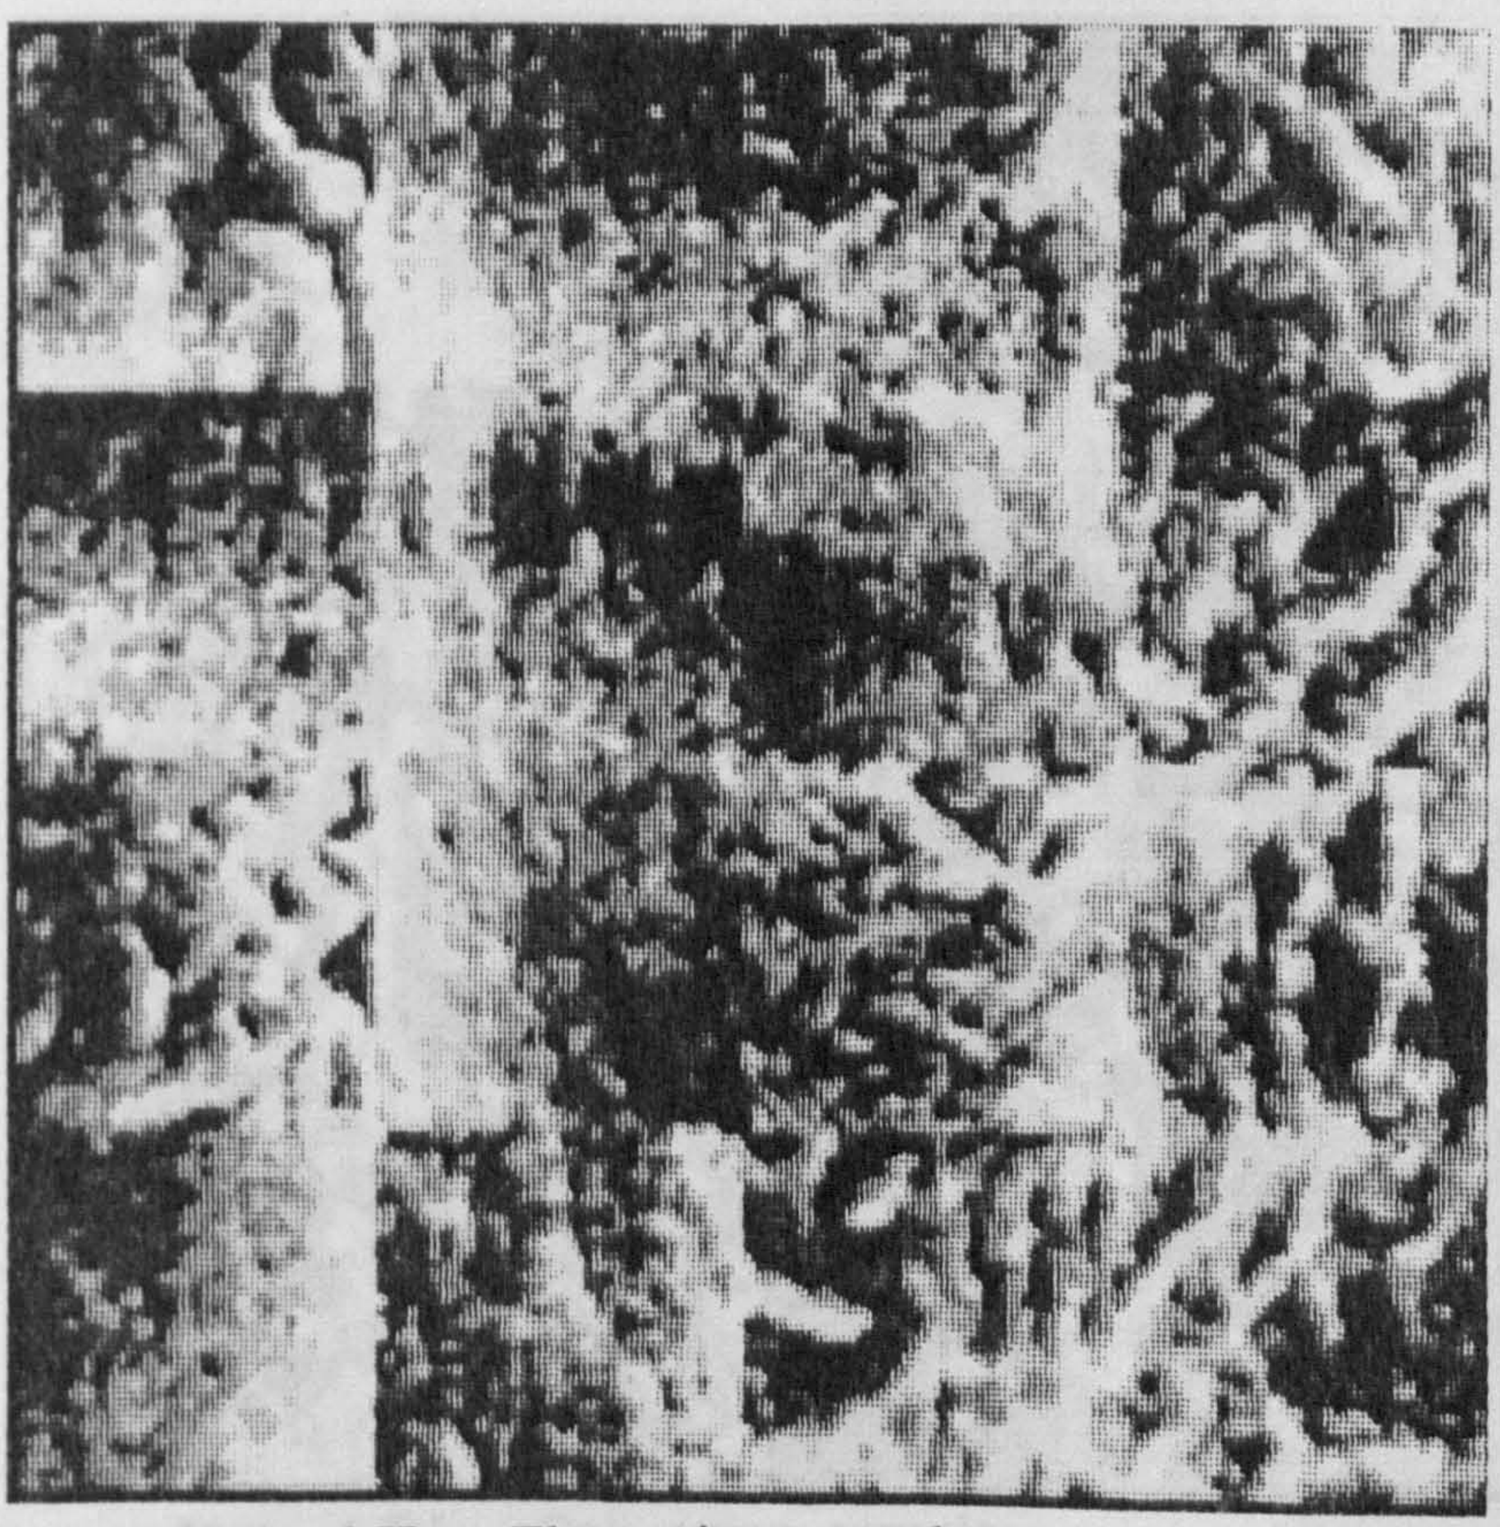

Plate 6-II. Elm exine samples.

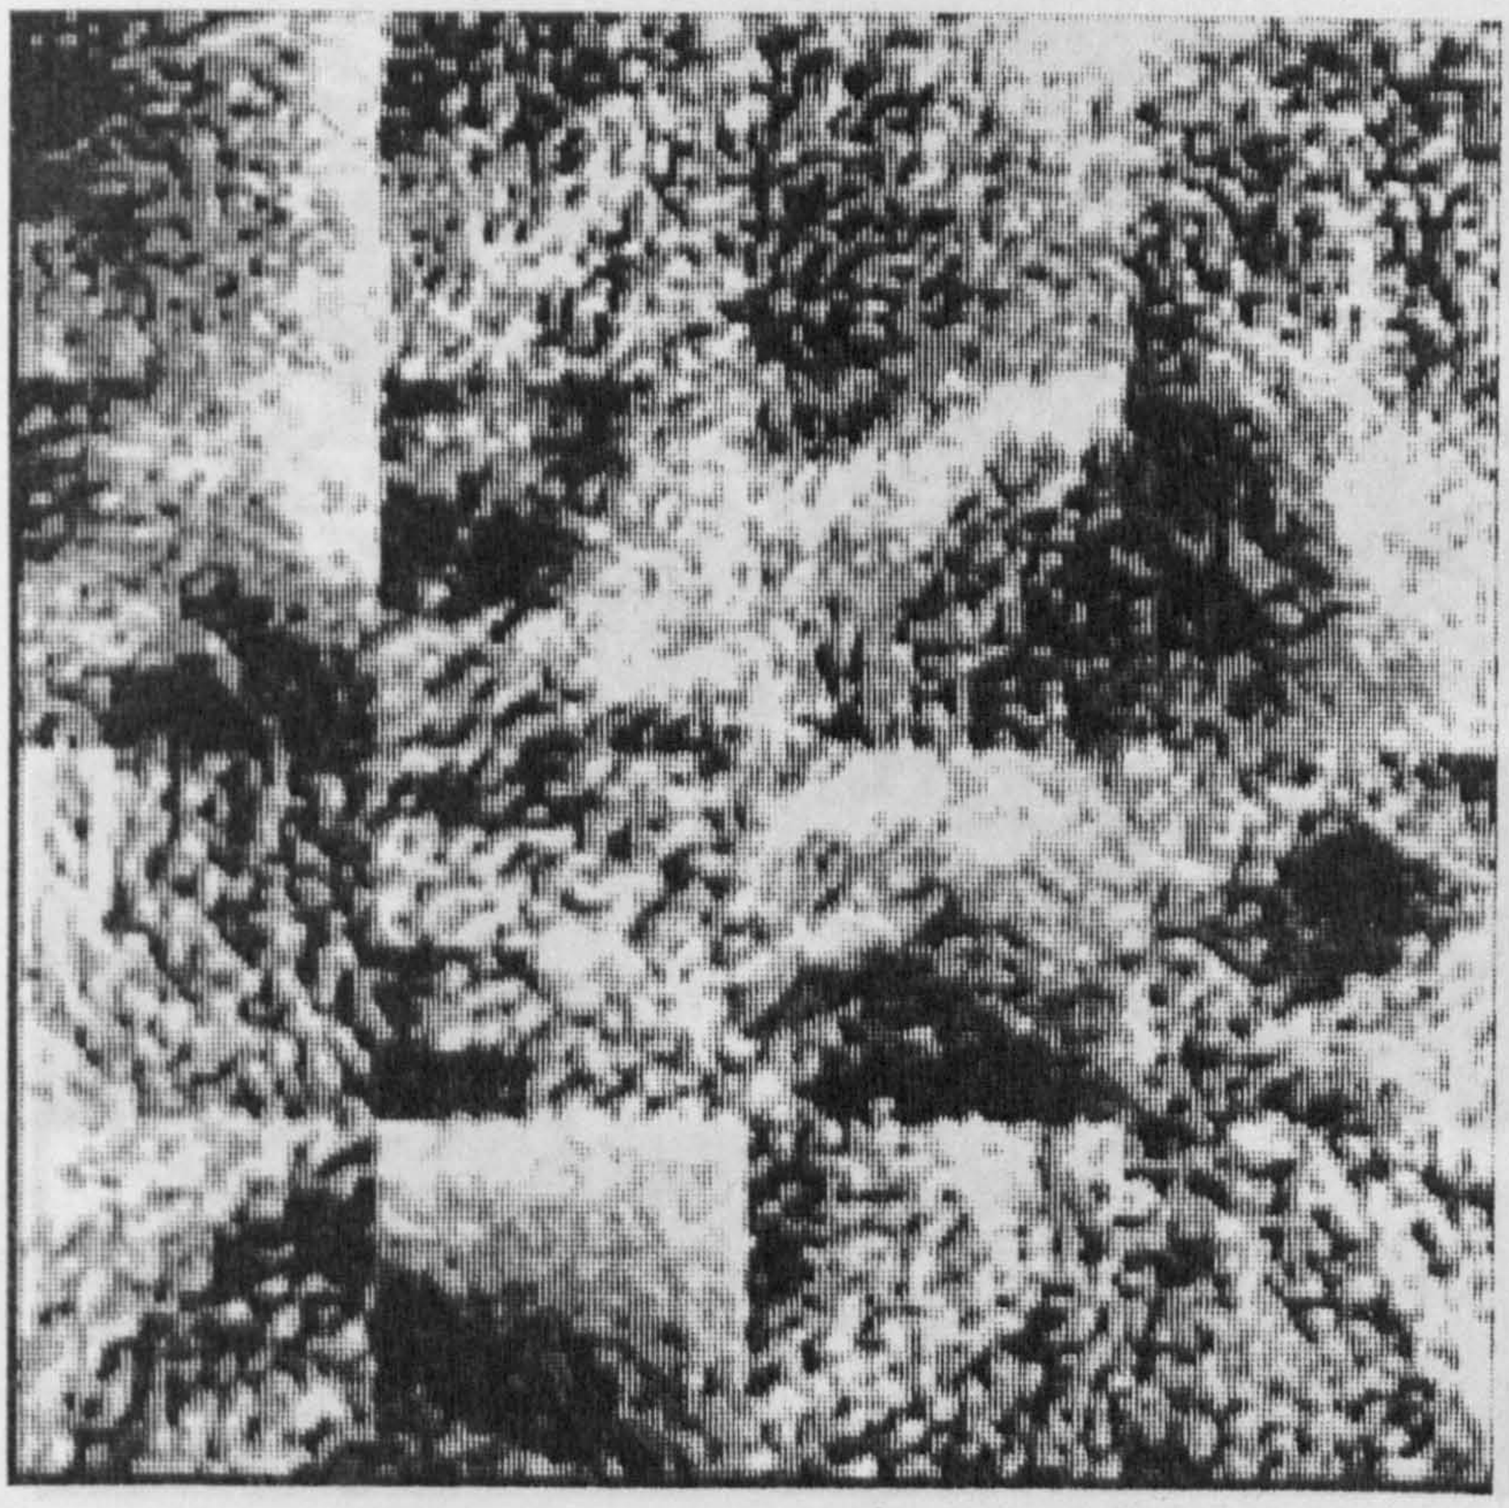

Plate 6-III. Rye grass exine samples.

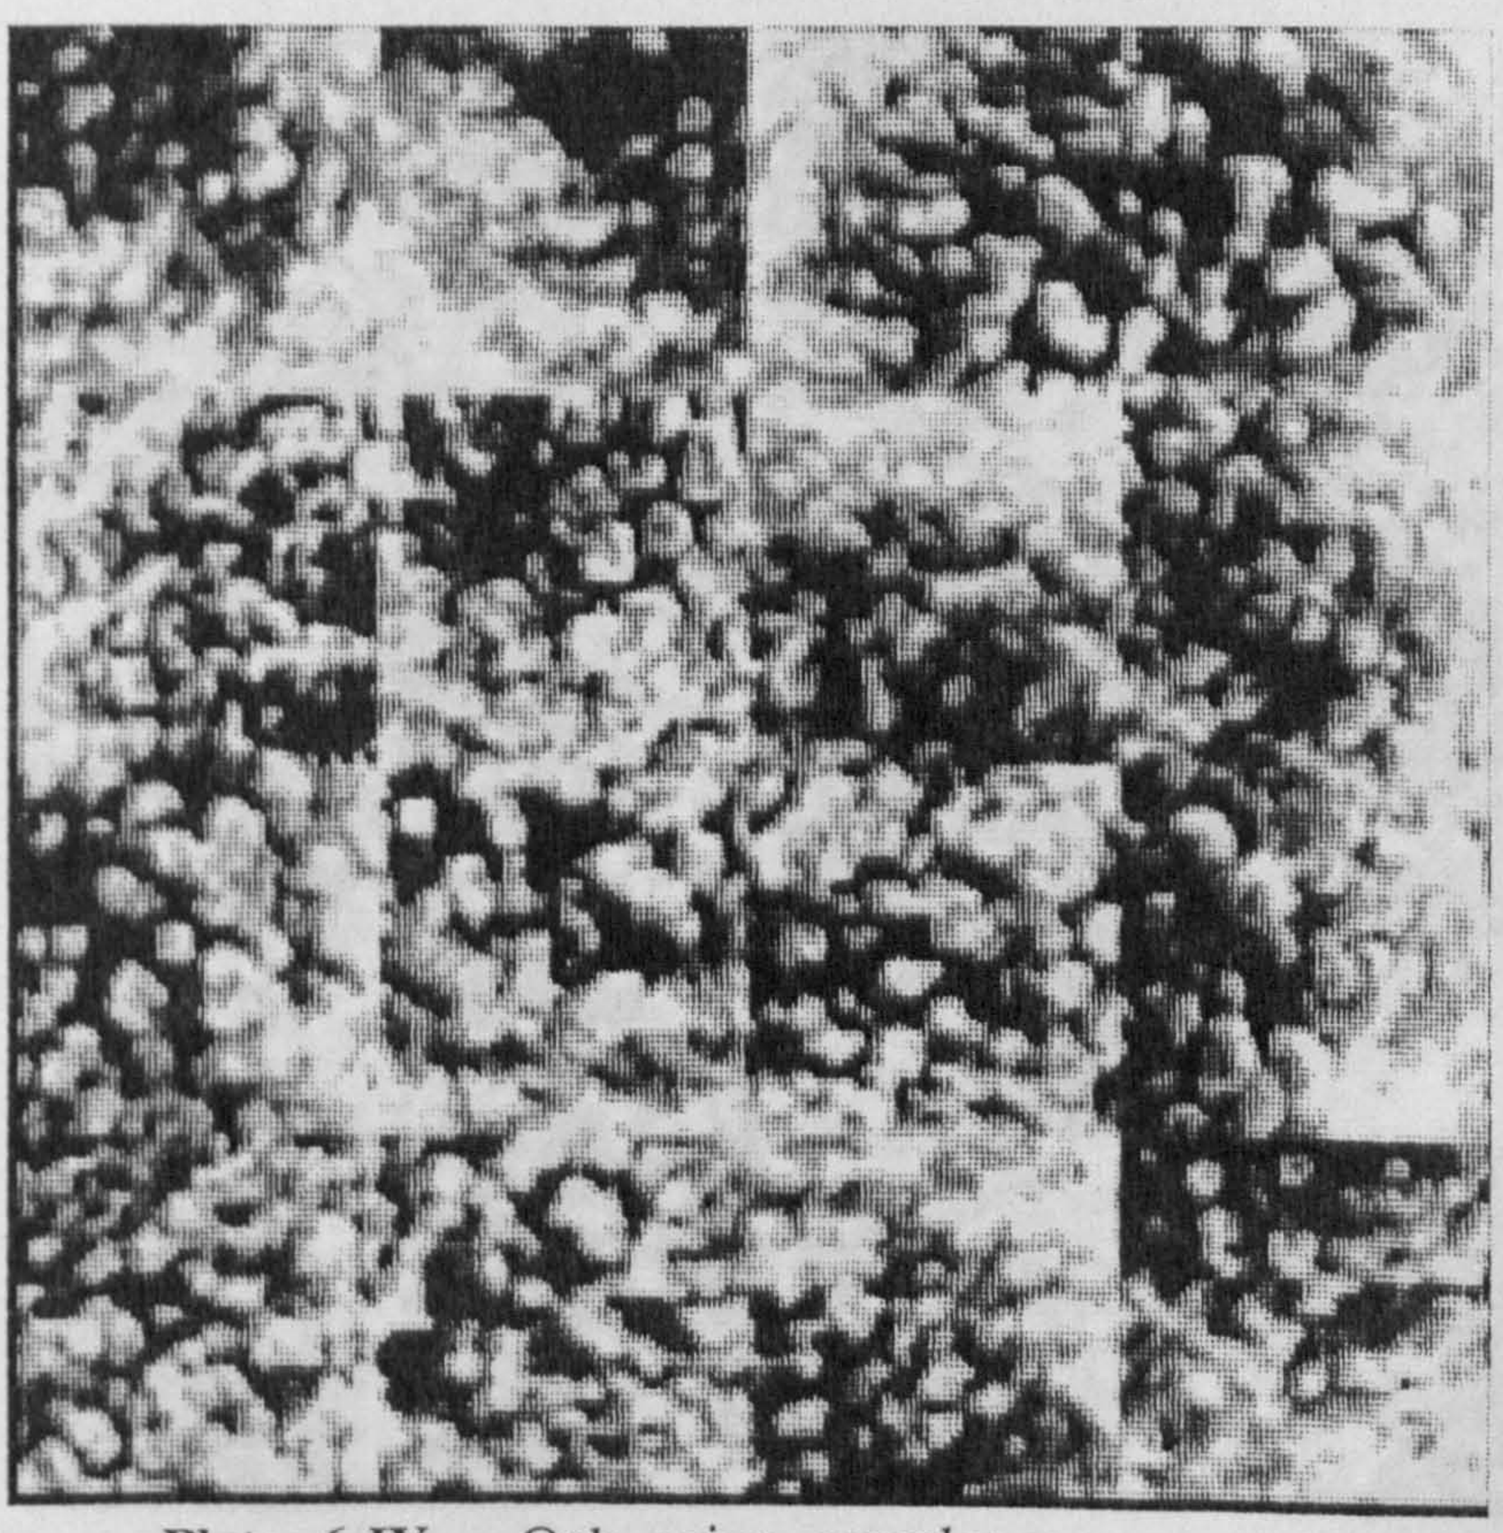

Plate 6-IV. Oak exine samples.

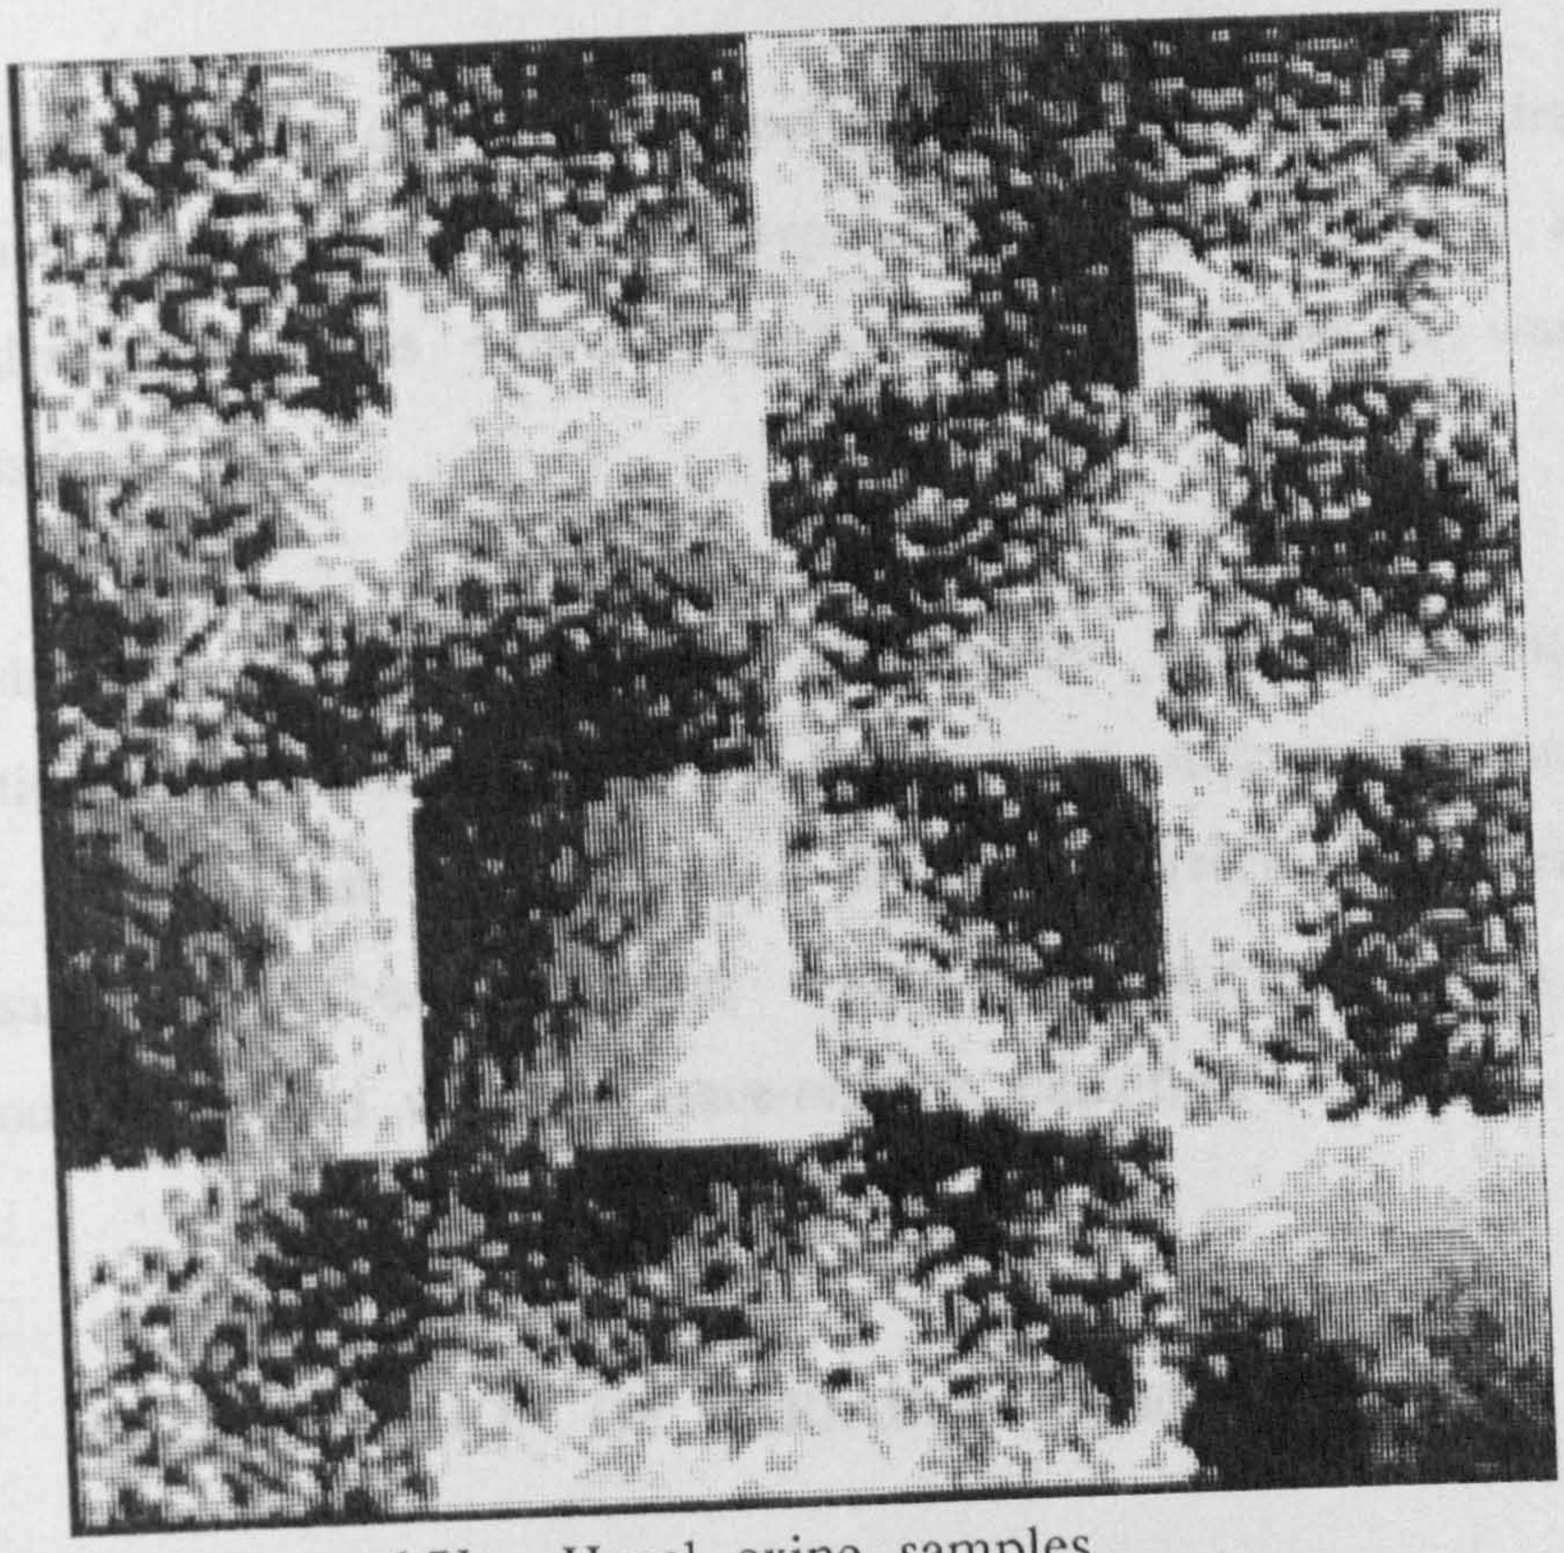

Plate 6-V. Hazel exine samples.

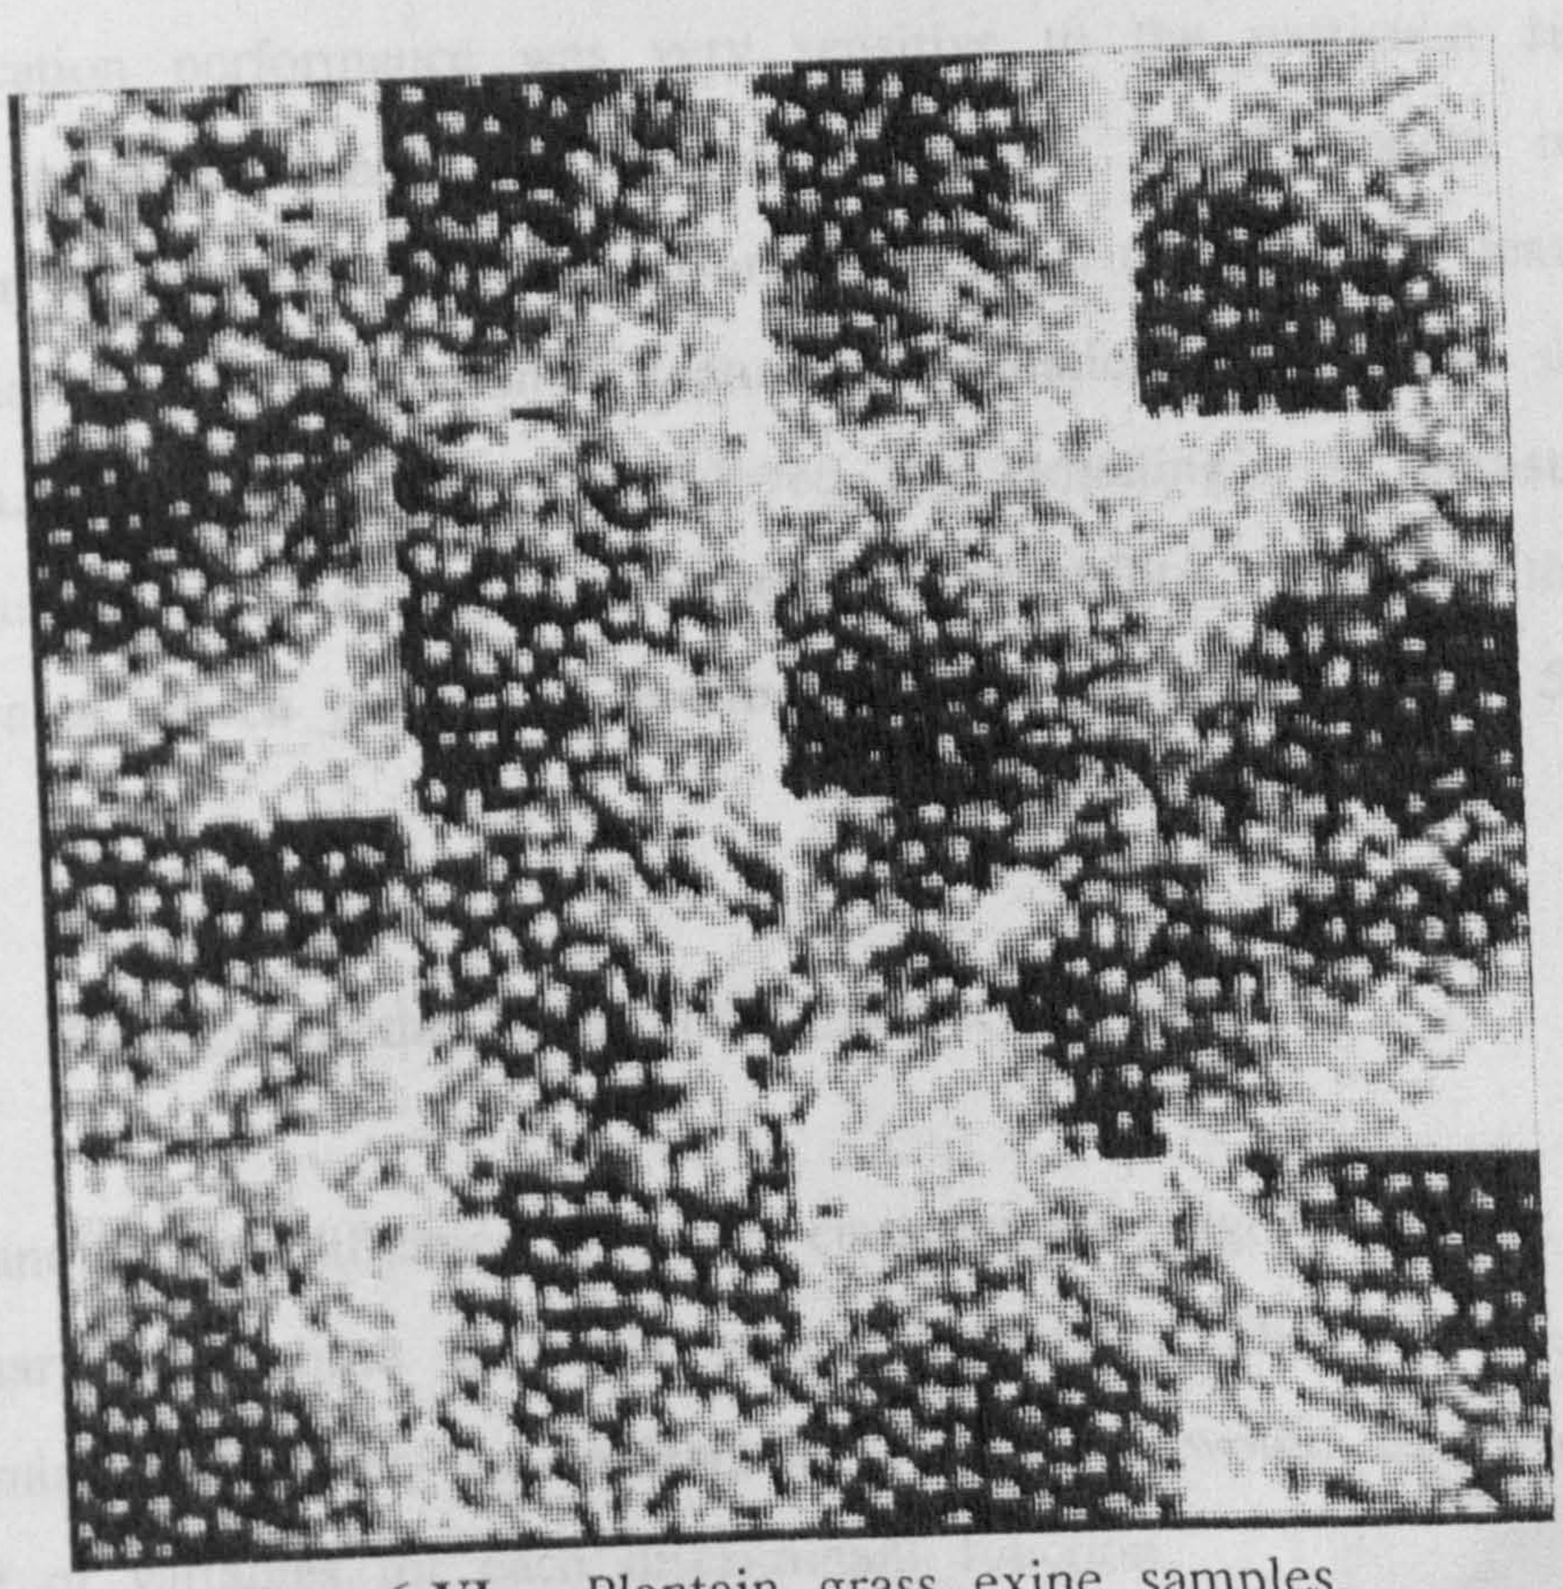

Plate 6-VI. Plantain grass exine samples.

reentered into the discriminant function constructed for the drawn class pair and the output from this taken as the final result. If a sample was drawn on votes between more than two classes it was left unclassified.

Applying this 'drawn-between-two' tie-breaker decision algorithm and repeating the classification process the average success rate increased to 44%, and the best pair of texture measures correctly classified 55% of the samples. These improved results were comparable to those previously obtained with the leave-one-out classifier.

6.4 INCORPORATION OF VARIABLE SELECTION

It had been established from the results obtained so far that the classification performance was very sensitive to the particular subset of texture measures used. A selection of a few 'good' texture measures could rival the classification performance of a much larger feature set. In order to identify optimal feature combinations a variable selection procedure was used. This employed the Hotelling's T^2 statistic as a multivariate measure of between-class separability, along with the accelerated search technique described in Chapter Five (section 5.7.5).

6.4.1 Redesigning the Linear Discriminant Classifier

Yet another modification to the classification algorithm was now necessary, this time a rather major alteration. The Fisher linear discriminant classifier was redesigned so that it would use the optimal subset of variables for each discriminant function.

Understandably the optimal features for discrimination varied depending on which particular class pair was under consideration. The variable selection program produced a data file in which the optimal subset of variables for discriminating each individual class pair were held. The classifier algorithm was modified to access this file and used only the selected features when constructing the linear discriminant function for each class pair combination. Similarly, when it came to classifying the unknown samples, the same file was accessed so that only the appropriate features were used in each discriminant function calculation. The leave-one-out scheme could still be retained in this more complex classifier design. The results obtained with the new scheme are presented below.

6.4.2 Improved Classification Results

A summary of classification success rates utilizing variable selection in a Fisher linear discriminant classifier, are displayed in Table 6.5. Even without employing the tie-breaker decision algorithm over 66% of the samples were correctly identified using two texture measures. This was a very substantial improvement over the result that had been obtained without feature selection (44%). If the tie-breaker decision was also employed a very impressive 93% of the 192 samples were correctly identified, again an enormous improvement.

When the decision algorithm was omitted a steady improvement in the results was observed as more texture measures were incorporated. However, when using the decision algorithm no substantial improvement was obtained with larger feature sets. It appeared that this new classifier design incorporating both variable selection and the tie-break

Number of Features Used in Classifier	Tie-breaker Decision:	
	No	Yes
2	66.7	92.7
3	75.5	87.0
4	78.7	90.6
5	81.8	92.7
6	80.7	93.0

Table 6.5 Classification results using feature selection, both with and without a decision taken on samples drawn between class pairs.

decision algorithm was performing about as well as possible with the available data. No advantage could be gained from employing a larger number of texture measures.

Evidently the incorporation of variable selection and the modified classifier design lead to a powerful classification scheme. The improvements produced by variable selection and a decision between drawn class pairs is summarized in Table 6.6 below. This shows the success rate that was obtained, using only two texture measures, under the various classification schemes available.

	Tie break off	Tie break on
Selection off	25%	44%
Selection on	67%	93%

Table 6.6 Summary of the improvements to classification success gained through modifications of the classifier design.

6.5 ANALYSIS BY TEXTURE MASKS

6.5.1 Texture Measures and Classification Procedure

The so called Laws' masks, described by Pietikainen et al. (1983), were used as an alternative texture analyzer to grey level co-occurrence. When these texture filters are convolved with the image each mask produces a single texture descriptor as an output. Five masks of 3x3 dimension, and eight masks of 5x5 dimension were available. Thus five or eight texture measures could be extracted from a texture sub-image.

The image database was unaltered from the 192 grey level equalized samples used in the previous section, as these were quite suitable for analysis with the spatial filters. Similarly, the variable selection and tie-breaker decision algorithms were unaltered from their previous application. The results using texture mask analysis are presented below.

6.5.2 Results with Laws' Masks

The results for the 3x3 masks, exhibited in Table 6.7, were very promising. Utilizing two selected features 75% of the samples were correctly identified. The error rate dropped as more features were deployed, and the best result of 85% was obtained with four texture descriptors. Although these success rates were not as good as the co-occurrence results it was possible to calculate these texture measures more quickly. They also suggested that the more sophisticated 5x5 masks may be able to rival the performance of co-occurrence analysis.

Features Used	3x3 masks	5x5 mask
2	75.0	84.4
3	81.8	96.4
4	84.9	96.4
5	84.4	96.4
6	-	96.4
7	-	96.4
8	-	95.8

Table 6.7 Classification results using 3x3 and 5x5 Laws' masks.

The 5x5 masks did indeed produced better results. A success rate of over 84% was achieved by two features, and a consistent 96.4% was maintained as more features were utilised in the classifier. Indeed these results, with the exception of the two feature subsets, exceeded the success rates of co-occurrence analysis. With five texture measures used, three of the samples in the database remained unclassified. If these are treated as 'rejects' then 98% of the samples were identified correctly. This is a very satisfactory outcome.

It would appear therefore that the Laws' mask texture filters are a powerful scheme for texture analysis. Certainly in these experiments the 5x5 masks produced better results than co-occurrence analysis, which itself performed admirably.

6.6 COMPARISON OF LAWS' MASKS AND CO-OCCURRENCE ANALYSIS

The confusion matrices displayed in Tables 6.8(a-c) show the classification results obtained by Laws' masks and co-occurrence analysis when using five texture measures. They confirm the superiority of the 5x5 masks over co-occurrence analysis.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**100.0**	0.0	0.0	0.0	0.0	0.0	0.0
El	9.3	**84.4**	6.3	0.0	0.0	0.0	0.0
Ry	6.3	0.0	**93.7**	0.0	0.0	0.0	0.0
Oa	0.0	0.0	0.0	**100.0**	0.0	0.0	0.0
Ha	6.3	0.0	0.0	0.0	**90.6**	0.0	3.1
Pl	9.3	0.0	0.0	0.0	0.0	**87.5**	3.1

Table 6.8a Confusion matrix for samples classified using co-occurrence texture measures. Figures are percentages.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**93.7**	0.0	0.0	0.0	0.0	0.0	6.3
El	6.3	**87.5**	3.1	0.0	0.0	0.0	3.1
Ry	9.3	3.1	**84.5**	0.0	0.0	0.0	3.1
Oa	18.7	0.0	3.1	**78.1**	0.0	0.0	0.0
Ha	3.1	3.1	6.3	0.0	**84.5**	0.0	3.1
Pl	15.6	0.0	6.3	0.0	0.0	**78.1**	0.0

Table 6.8b Confusion matrix for samples classified using Laws' 3x3 texture measures. Figures are percentages.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**100.0**	0.0	0.0	0.0	0.0	0.0	0.0
El	6.3	**90.6**	3.1	0.0	0.0	0.0	0.0
Ry	0.0	0.0	**100.0**	0.0	0.0	0.0	0.0
Oa	3.1	0.0	0.0	**96.9**	0.0	0.0	0.0
Ha	0.0	0.0	0.0	0.0	**90.7**	0.0	9.3
Pl	0.0	0.0	0.0	0.0	0.0	**100.0**	0.0

Table 6.8c Confusion matrix for samples classified using Laws' 5x5 texture measures. Figures are percentages.

The variable selection procedure supplies not only an optimal subset of texture measures but also a corresponding T^2 statistic. This measure of multivariate class separability may give some insight into the performances of the different texture analyzers. In particular we may investigate the reason why Laws' masks do less well than co-occurrence when only two texture measures are employed, but become superior with five feature classification.

Table 6.9 presents the T^2 values corresponding to the optimal subsets selected from Laws' mask and co-occurrence texture measures. These are given for subsets consisting of two and five features.

Considering the T^2 values of two feature subsets between Pine and all other classes it is noticeable that whilst the figures are generally high, the Laws' values are substantially smaller than the co-occurrence. With subsets of five features this difference is greatly reduced, the Laws' features having all but caught up with the co-occurrence values. On the other hand, Laws' subsets of two features have greater distinguishing power than co-occurrence measures for class pairs containing Elm, but this difference is also reduced for five feature subsets. For class pairs containing Rye grass the Laws' T^2 values show a much bigger improvement between the two and five feature subsets than co-occurrence. Similarly the Laws' subsets appear to show a greater improvement for class pairs containing Oak, Hazel, and Plantain.

The general impression is that co-occurrence subsets of two features have greater overall discriminating power than the Laws' subsets. However, when five features are used the Laws' subsets show the greatest improvement and become superior to the co-occurrence measures. In order to back up this contention we may consider the median or the

	Pi	El	Ry	Oa	Ha
El	228 111				
Ry	896 234	53 180			
Oa	206 95	63 150	292 94		
Ha	783 117	23 54	33 38	167 248	
Pl	783 117	84 146	80 60	112 82	42 59

	Laws' Masks	Co-occurrence
Sum of Ranks	229	236

(a) T^2 values for two variable subsets.

	Pi	El	Ry	Oa	Ha
El	356 382				
Ry	1576 1092	94 310			
Oa	318 374	196 190	418 426		
Ha	914 887	112 195	75 172	253 262	
Pl	1060 668	169 403	168 437	786 196	95 89

	Laws' Masks	Co-occurrence
Sum of Ranks	250.5	214.5

(b) T^2 values for five variable subsets.

Table 6.9 Comparison of T^2 values from optimal subsets of co-occurrence and Laws' masks texture measures. Upper figures are derived from co-occurrence measures and lower figures from Laws' 5x5 masks.

sum of ranks of the T^2 figures. The mean would be inappropriate measure due to the highly skewed distribution of these data. Table 6.9 also displays the sum of ranks. It can be seen that the co-occurrence values have a slightly larger sum of ranks for two feature subsets, indicating a slim superiority in discriminating power. This situation is reversed for the five feature values, and furthermore the Laws' measures hold a much larger superiority. A lower degree of correlation between the variables is the probable explaination as to why the Laws' subsets become more powerful than co-occurrence subsets as further texture measures are incorporated.

6.7 COMBINED CO-OCCURRENCE AND TEXTURE MASKS

It was felt that combining Laws' mask and co-occurrence texture measures may produce subsets with even greater discriminating power. To examine this idea the Laws' texture measures were combined with eight of the co-occurrence measures, selected somewhat arbitrarily, and entered into the variable selection procedure. The full set of co-occurrence measures were not used because the variable selection procedure would have become very slow with such a large number of variables to process.

In general the T^2 values for a given class pair were greater than those obtained when the Laws' and co-occurrence measures were used alone. However, a few of the two feature subsets did not achieve as high a degree of separation as the co-occurrence subsets, presumably because a critical co-occurrence measure had been omitted in the arbitrary selection. It was also noticeable that with subsets of five or more variables the Laws' texture measures were more frequently selected than

co-occurrence measures. This is perhaps another indication of their greater discriminating power.

The results using the combined measures are presented in Table 6.10 below. The 88% achieved with two features was better than the Laws' mask features managed, but not as good as the co-occurrence features when they were used exclusively. This was to be expected considering the T^2 values described above. With four or more features the 98.4% success rate was an improvement over the figures for the separate use of either texture analyzer. In fact this success rate meant that only three samples were incorrectly identified.

Features Used	% Correct Classification
2	88.0
3	95.8
4	98.4
5	97.9
6	98.4
7	98.4
8	98.4

Table 6.10 Classification success from combined co-occurrence and Laws' mask texture measures.

It would appear, therefore, that combining measures from the two texture analyzers can produce feature subsets with greater discriminating power than those obtained from Laws' mask or co-occurrence measures alone. This is clearly reflected by the greater classification success rates obtained with combined measures.

6.8 ADDITION OF A STRUCTURAL TEXTURE ANALYZER

The classification success rates obtained in the previous section were considered to be very encouraging. A palynologist would expect misclassification errors of up to 5% to occur during manual identification. However, the palynologist would be attempting to classify a greater number of pollen taxa, and would generally be concerned with fossilized pollen which are likely to prove a greater challenge for automated recognition. An automated system may also be required to differentiate between more than six taxa, and this might increase the error rates. Furthermore, it may also be necessary to differentiate between different taxa for which the statistical texture analyzers described above would be less suitable. The samples used up until now had all been taken at the same magnification and only modified by the histogram equalization technique. Altering the magnification and applying more preprocessing operations may also have a significant affect on the ability to classify correctly the exine samples.

For these reasons it was decided to continue experiments into classification by texture analysis. A structural texture analyzer was added to the statistical methods which required the samples to be processed differently. The effects of these changes are described below.

In this section some consideration was also given as to how the complexity of the system might be reduced. The variable selection and classification procedures used so far were compared to some simpler methods. The results of these comparisons are also discussed below.

6.8.1 Details of the Database

In order to make use of the structural texture analyzer based upon matched edge pairs (Pietikainen and Rosenfeld, 1982) it was necessary to generate a new database with images whose characteristics were more suitable for this analysis scheme. In particular, it was necessary that the textural primitives (spots, streaks, blobs, etc.) were fairly large, and that noise and consequent false edges in the image were minimized.

A fresh set of exine samples was selected from new SEM photomicrographs of the same six pollen taxa. Sub-scenes of 64x64 pixels were extracted, but the magnification of the samples was increased. Of course, this meant that the actual portion of the pollen grain covered by a sub-scene was considerably less than with the previous samples. The change in magnification was achieved at the digitization stage by altering the height of the camera on its stand. Sixteen samples were collected from each pollen taxon, producing a database with ninety-six samples in total. The new database is illustrated in Plates 6-VII to 6-XII.

6.8.2 Preprocessing of Exine Samples

The new samples were first processed with a 3x3 median filter. This was applied twice in order to greatly suppress noise levels in the image, and to eliminate very weak edges. However, it was hoped that the strong edges associated with texture primitive boundaries were not affected by the process and so would remain sharply defined. The median filter was selected in preference to a simpler mean filter in order to try and preserve these stronger edges. Using the median filter

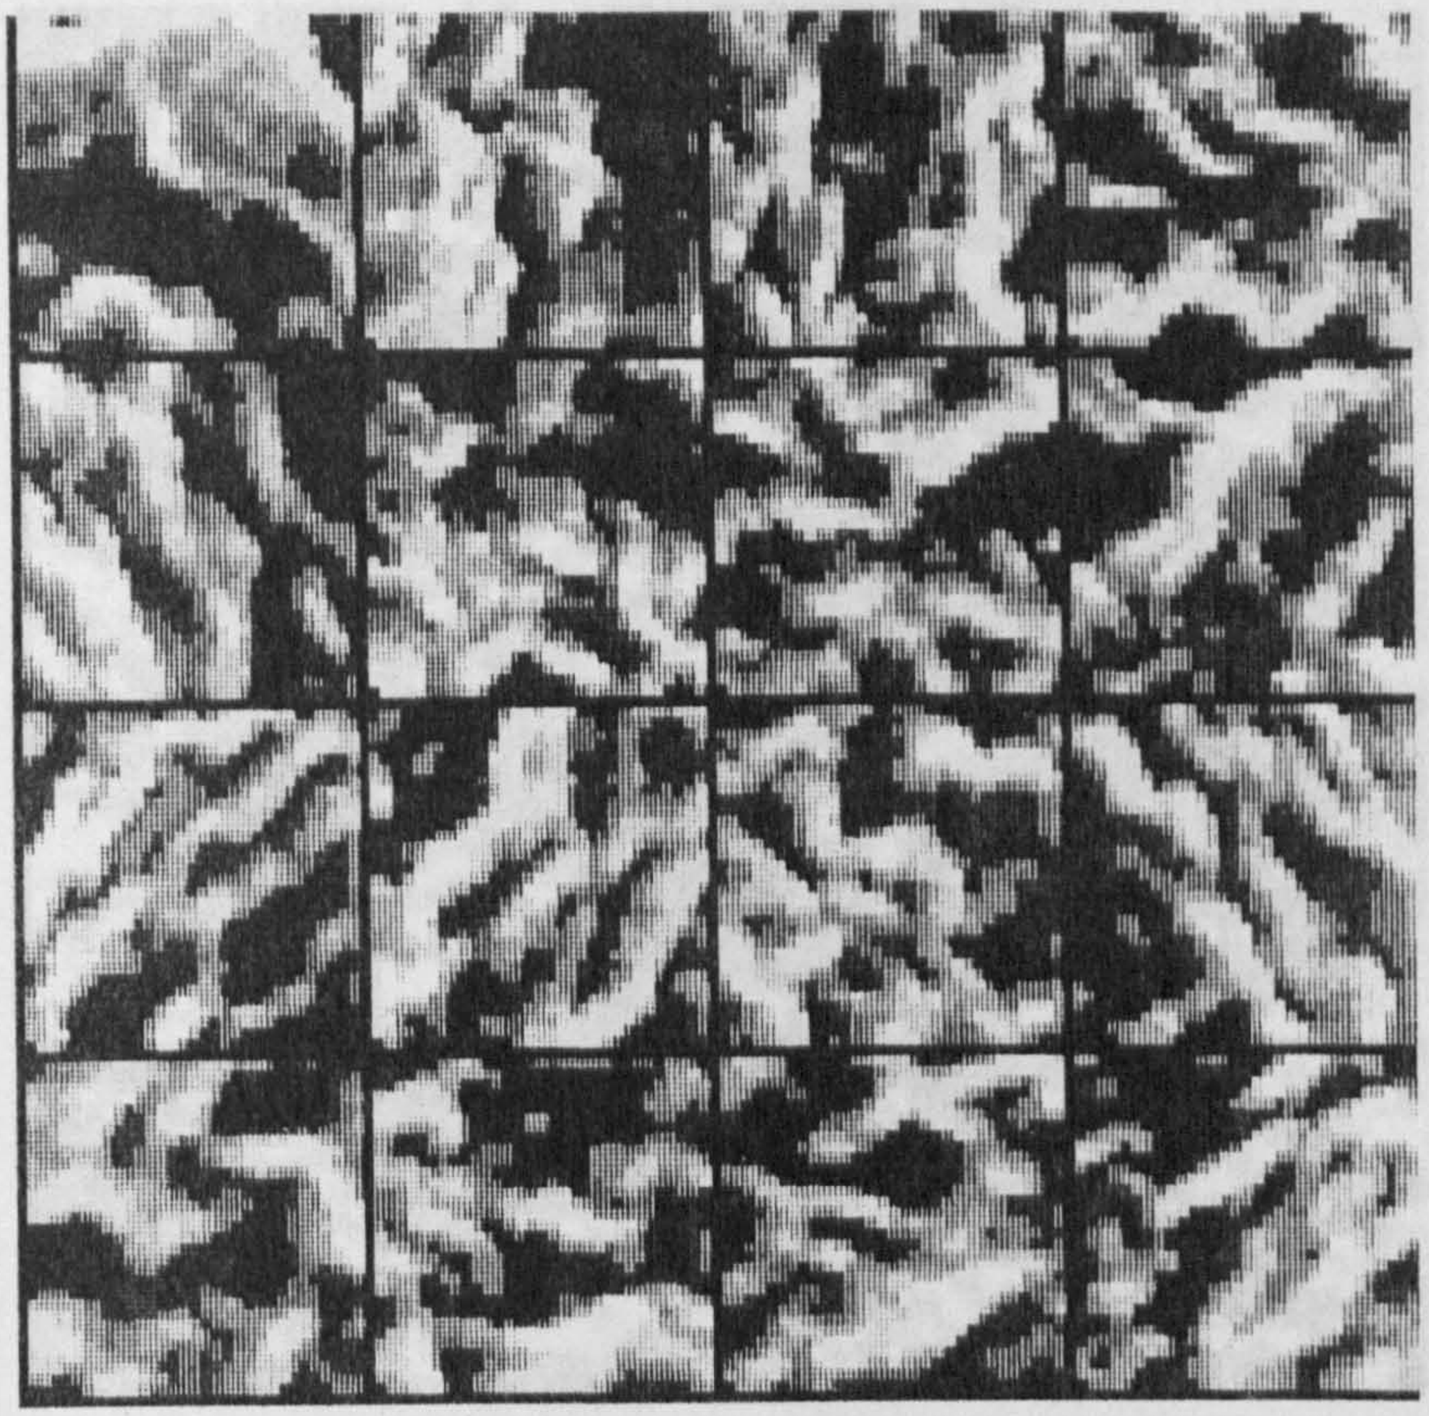

Plate 6-VII. Pine exine samples.

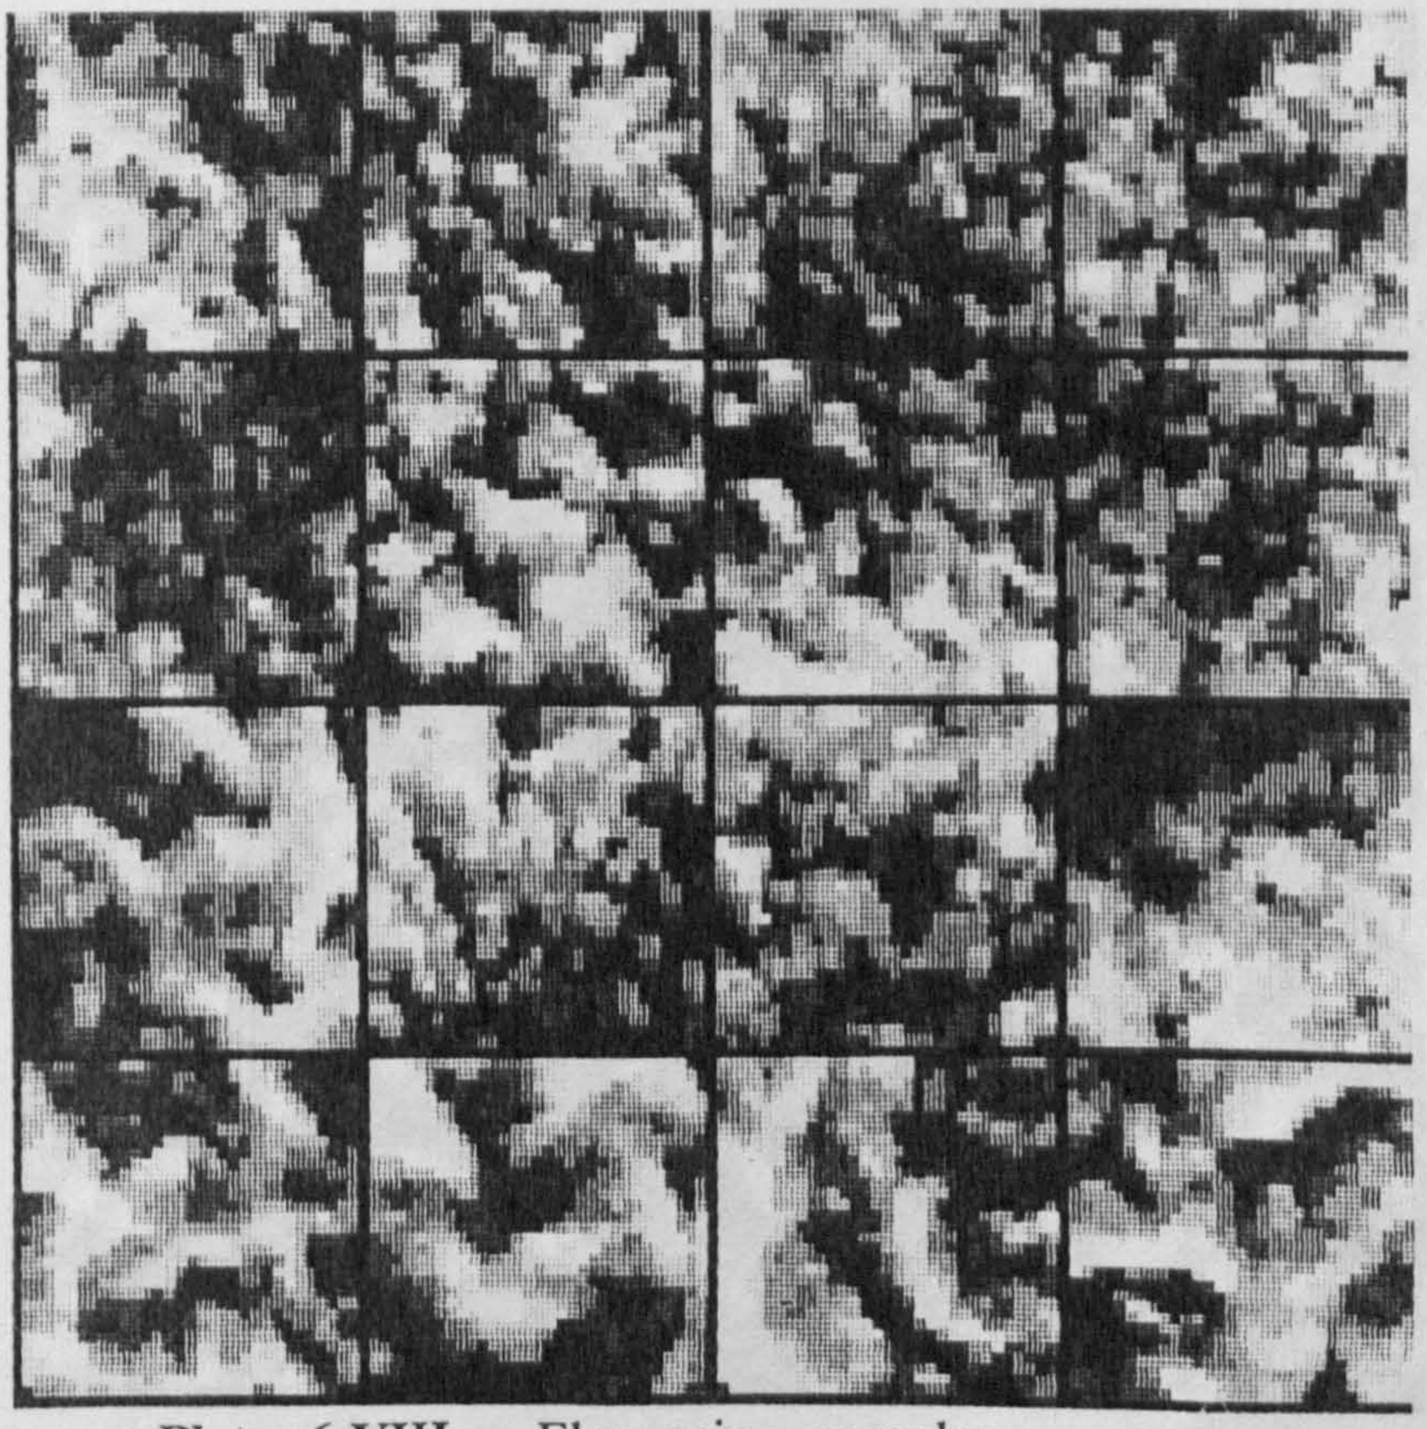

Plate 6-VIII. Elm exine samples.

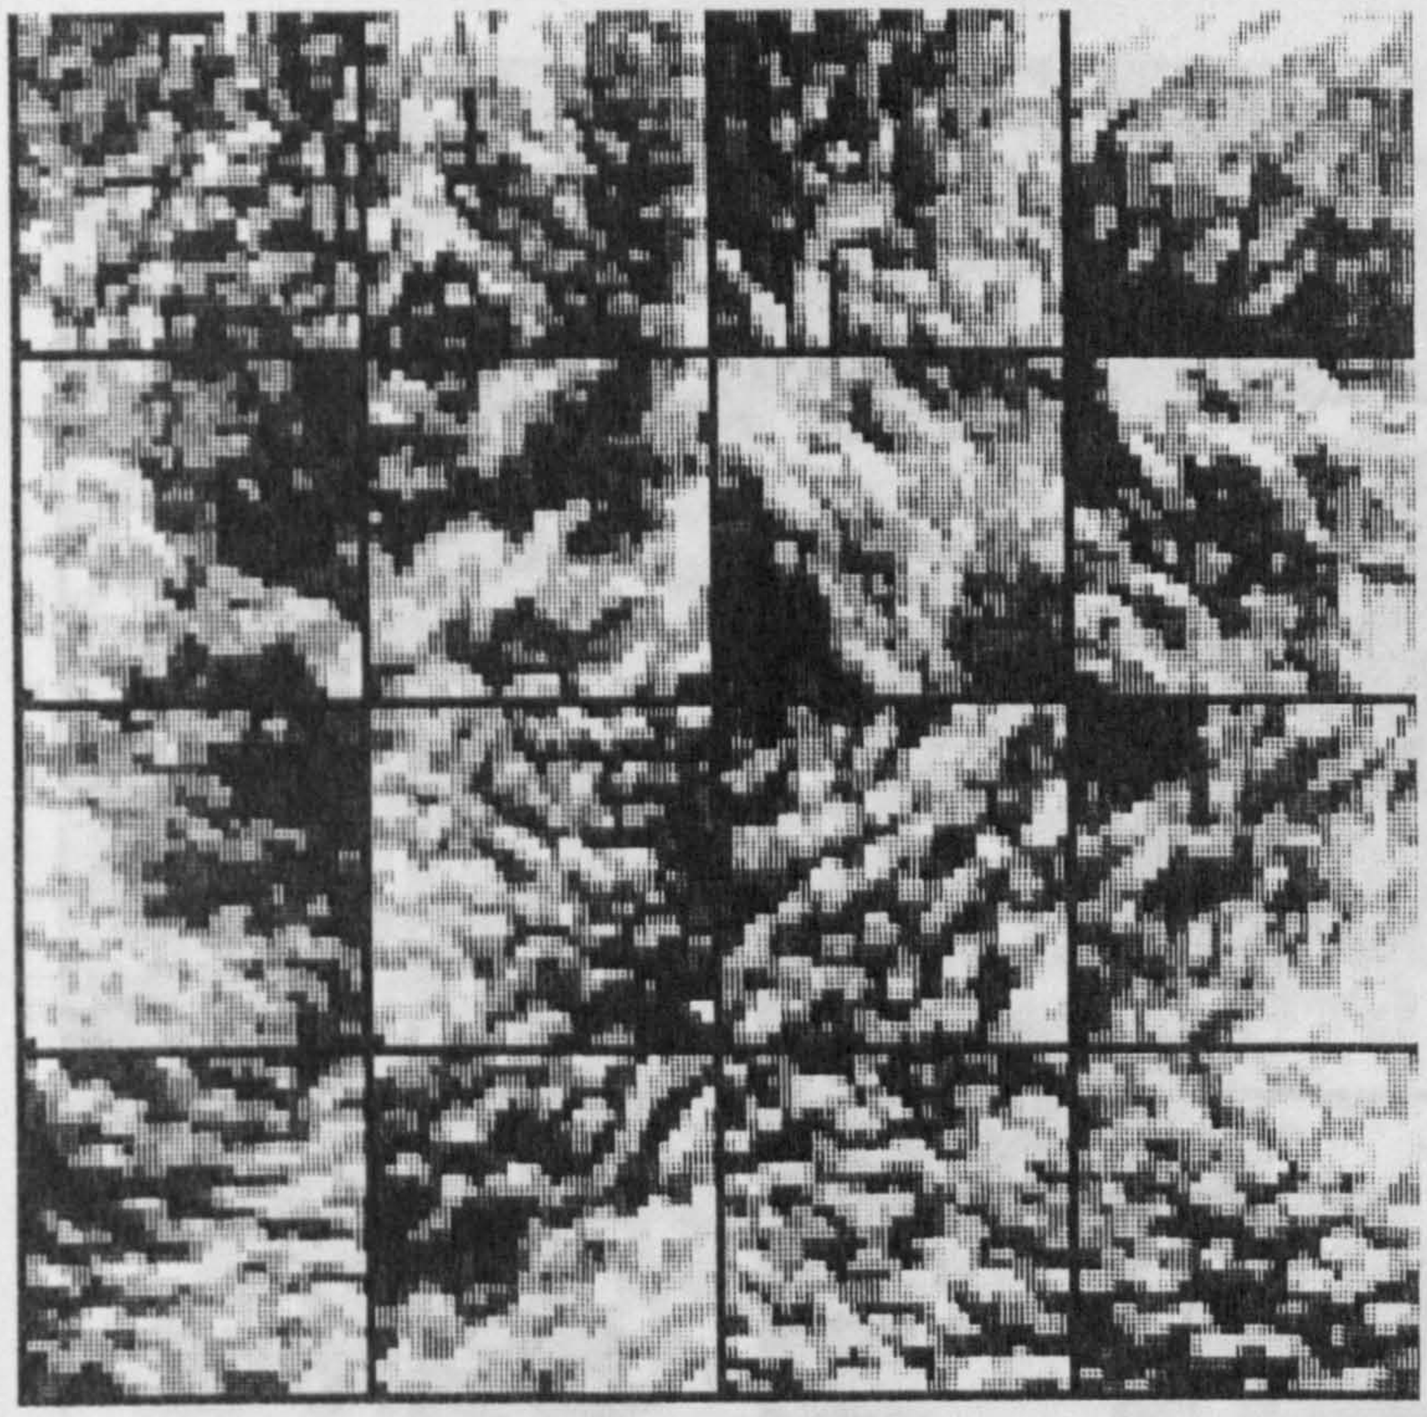

Plate 6-IX. Rye grass exine samples.

Plate 6-X. Oak exine samples.

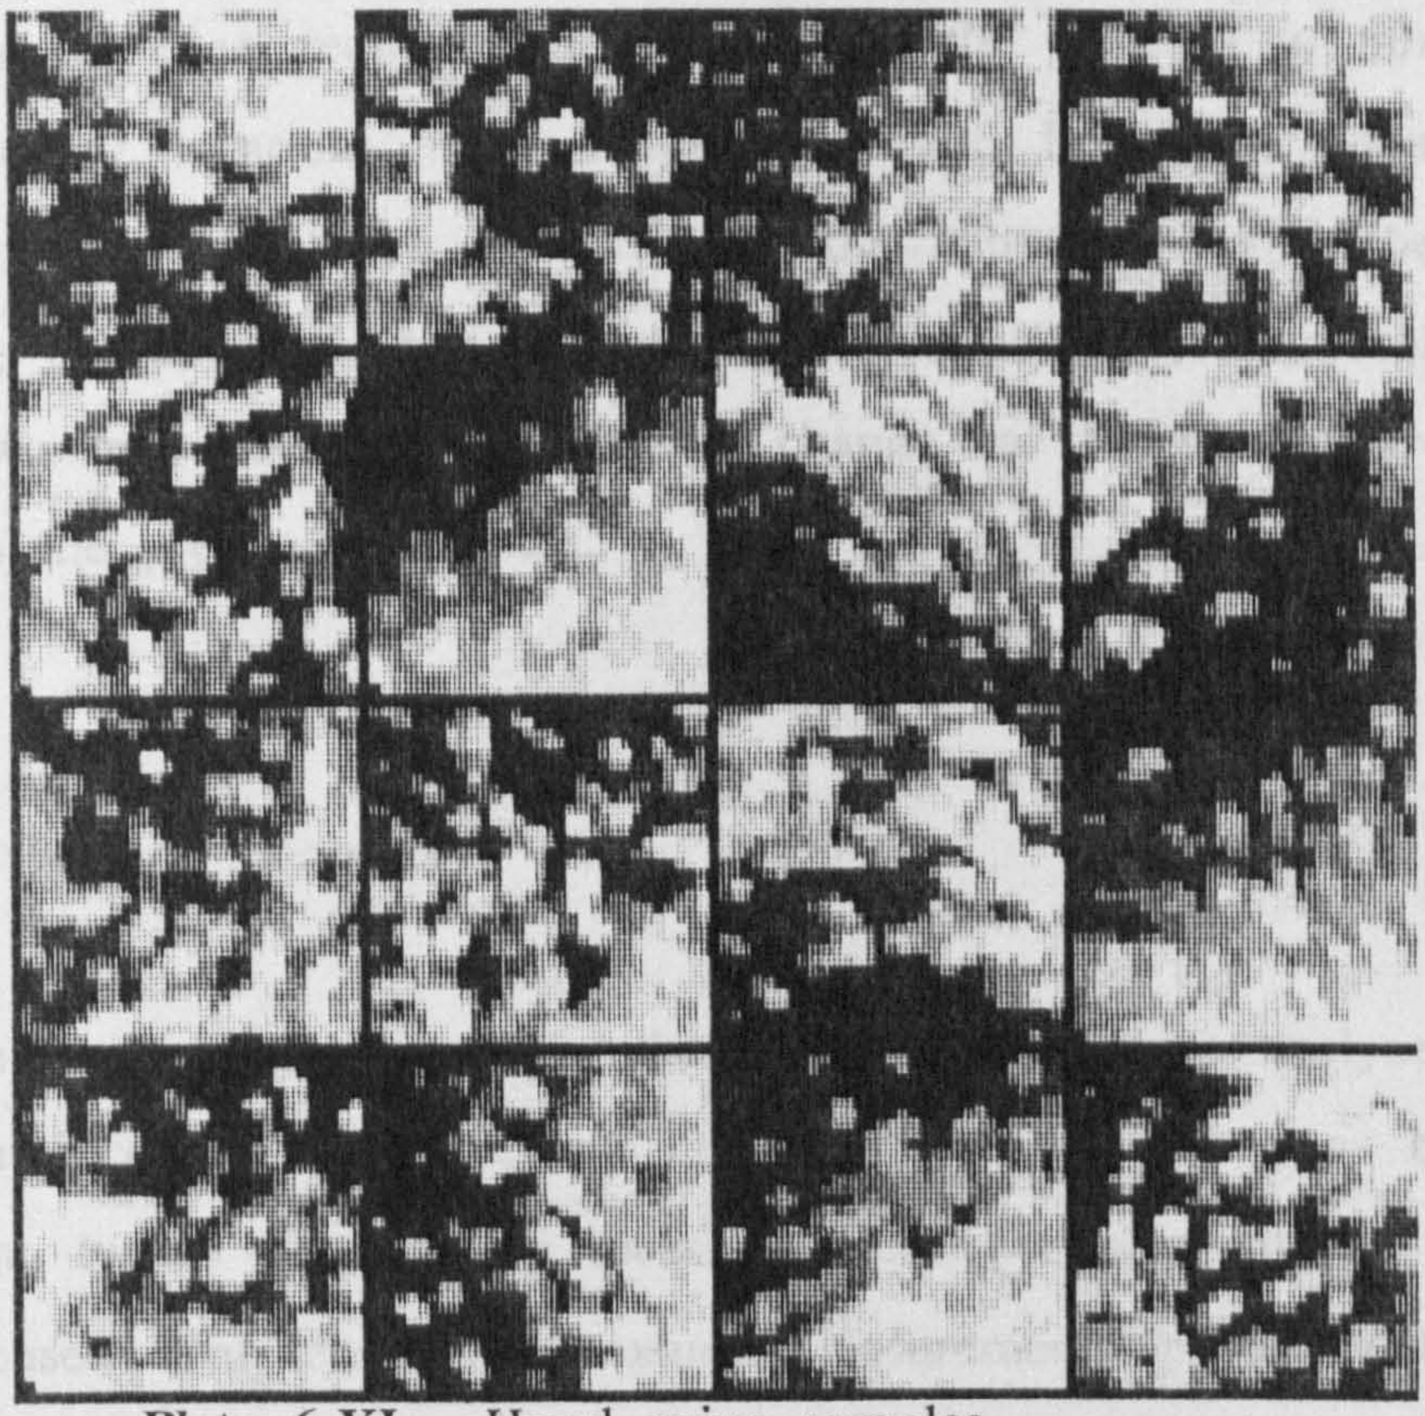

Plate 6-XI. Hazel exine samples.

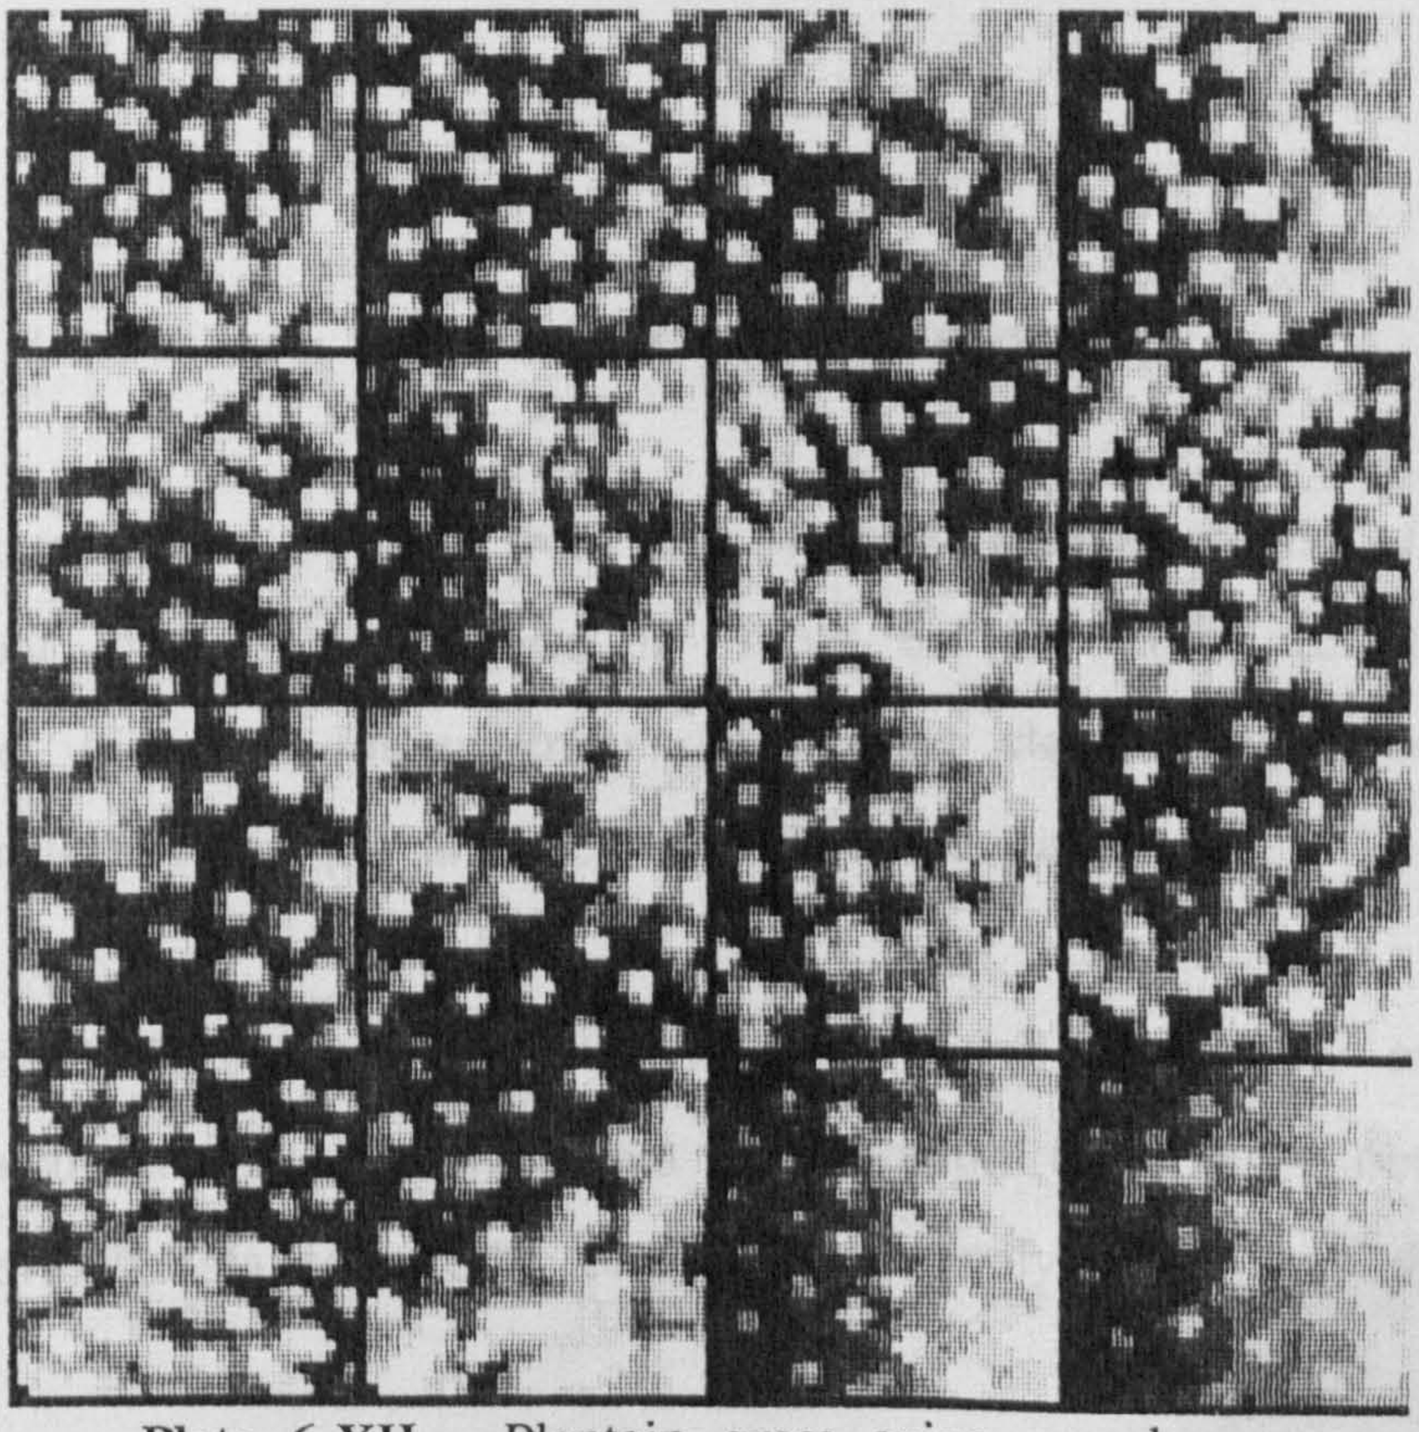

Plate 6-XII. Plantain grass exine samples.

operator caused a reduction in the size of the samples to 60x60 pixels as boundary locations were discarded. The histogram equalization technique was employed to normalize the sub-scenes after the median filtering operation. The samples were reduced in tonal resolution to thirty-two grey levels. This was a change to the sixteen grey levels used in the previous databases.

6.8.3 Texture Measures

Texture measurements were produced by three texture analyzers. The co-occurrence and Laws' texture measures were calculated as before. In addition, twelve new texture measures were produced by the structural edge based analyzer. The texture measurements produced by each analyzer were first used independently to classify the new database samples. Later they were combined in an attempt to improve classification results.

6.8.4 Results from Individual Texture Analyzers

The classification results from the texture measurements of each individual analyzer are shown in Table 6.11. These results were obtained using the Fisher linear discriminant classifier incorporating the leave-one-out technique and tie-breaker decision algorithm.

It is clear that the success rate of both the co-occurrence and Laws' mask analyzers show a significant deterioration from those obtained with the previous data set. The co-occurrence analyzer displays the smallest

Features Used	Co-occurrence	Laws Mask	Edge Pair
2	73	66	72
3	72	72	74
4	81	75	78
5	85	81	75
6	87	78	77
7	83	76	75
8	82	54	73

Table 6.11 Results from using the texture analysis schemes individually. Figures are percent correctly classified.

deterioration, with success rates averaging just over 80%, while the Laws' mask analyzer averages a little over 70%.

This deterioration in performance was expected due to the larger magnification of the new samples. It has been stated previously that macro textures are less appropriately analysed by grey level co-occurrence (Davis et al., 1979) since the derived statistics tend to reflect intensity transitions within the texture primitives rather than the large scale structural properties of the texture (Shen, 1980).

The same argument applies to the Laws' mask measures. The texture primitives in the newer samples were now much larger than the dimensions of the masks. Therefore, the masks would have to be increased in size if they were to respond to the local image properties that they were originally designed to detect.

The structural edge pair analyzer produced results that fall midway between those discussed above. They average around 75% success rate, but show less sensitivity to the number of texture features used.

Analysing the T^2 values for two feature subsets it was found that the co-occurrence measures produced the highest figures for most of the

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**87.4**	6.3	0.0	0.0	0.0	0.0	6.3
El	6.3	**81.1**	6.3	6.3	0.0	0.0	0.0
Ry	0.0	12.6	**87.4**	0.0	0.0	0.0	0.0
Oa	12.6	6.3	0.0	**74.8**	0.0	0.0	6.3
Ha	0.0	0.0	6.3	0.0	**56.1**	31.3	6.3
Pl	0.0	0.0	6.3	6.3	25.0	**56.1**	6.3

Table 6.12a Confusion matrix for classification by two co-occurrence texture measures.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**87.4**	0.0	0.0	6.3	0.0	0.0	6.3
El	6.3	**43.6**	12.5	6.3	12.5	12.5	6.3
Ry	0.0	6.3	**87.4**	0.0	6.3	0.0	0.0
Oa	0.0	6.3	0.0	**68.7**	0.0	25.0	0.0
Ha	0.0	6.3	0.0	43.6	**37.5**	0.0	6.3
Pl	0.0	12.5	6.3	0.0	6.3	**68.6**	6.3

Table 6.12b Confusion matrix for classification by two Laws' 5x5 mask texture measures.

Actual Class	Classified As: Pi	El	Ry	Oa	Ha	Pl	Uc
Pi	**56.1**	6.3	0.0	12.5	0.0	6.3	18.8
El	18.8	**81.2**	0.0	0.0	0.0	0.0	0.0
Ry	0.0	6.3	**81.2**	0.0	12.5	0.0	0.0
Oa	6.3	6.3	0.0	**68.6**	0.0	0.0	18.8
Ha	6.3	0.0	6.3	0.0	**68.6**	12.5	6.3
Pl	6.3	0.0	0.0	0.0	12.5	**81.2**	0.0

Table 6.12c Confusion matrix for classification by two edge pair texture measures.

class pairs. The edge pair measures produced far superior values for class pairs that included Plantain and closely rivalled the co-occurrence values for pairs that included Hazel. The Laws' mask features produced marginally better values for pairs that included Rye grass. If the confusion matrices shown in Table 6.12 are studied carefully it is seen that this pattern is closely mirrored in the results. Whilst the co-occurrence features yield the best overall results the edge pair features are much more successful in classifying the Hazel and Plantain samples. We might hope, therefore, that combining the co-occurrence and edge pair measures will yield greater classification success.

6.8.5 Comparison with Euclidean Distance Classifier.

It was decided that it would be useful at this stage to compare the performance of the Fisher linear discriminant classifier to the much simpler Euclidean distance classifier. We might expect the performance of the Fisher classifier to be better, but unless the margin is of sufficient magnitude it may be pertinent to use the simpler and quicker Euclidean classification scheme.

Table 6.13 displays the results obtained in a Euclidean distance classifier which incorporated the leave-one-out method. The performance was considerably poorer than the Fisher classifier, despite the fact that all the texture measures were employed during the classification. This proved unequivocally that the Euclidean classifier could not produce results of a sufficiently high standard. Therefore, the Fisher classifier is obviously to be preferred, even though it is substantially more complex and slower to compute.

Co-occurrence	Laws Mask	Edge Pair
49%	43%	60%

Table 6.13 Results obtained from a Euclidean distance classifier.

6.8.6 Comparison of Variable Selection by Accelerated Search and Sequential Backward Elimination

Another way in which the complexity of the system could be reduced would be to employ a variable selection procedure based on sequential backward elimination or sequential forward addition. rather than the accelerated search technique utilised so far.

To test how the classification performance might be affected by a simpler technique the texture measurement data were entered into a sequential backward elimination procedure. In this process we begin with the complete set of texture features and delete, one at a time, the variable which causes the smallest deterioration in the separability index. It is a similar technique to the accelerated search except that the relationships between discarded variables are ignored. It has the advantage, however, of being simpler and quicker to compute.

Table 6.14 shows the difference between the percentage classification results obtained from subsets selected by backward elimination and the accelerated search. Analysis of the T^2 values indicated that the separation levels of the subsets selected by backward elimination were generally a little poorer. The classification results remained comparable, however, with the co-occurrence measures, but were rather poorer with the Laws' masks and edge pair texture measures.

Features Used	Co-occurrence	Laws Mask	Edge Pair
2	+ 4%	-12%	- 3%
3	- 4%	- 5%	0%
4	+ 3%	+ 1%	- 2%
5	0%	0%	+ 3%
6	- 1%	0%	0%
7	- 1%	0%	0%
8	+ 2%	0%	0%

Table 6.14 Relative performance of the classifier after using variables selected by backward deletion rather than the accelerated search.

It is very noticeable that in some cases the performance improved with the use of the suboptimal feature sets. The final classification result for any given sample depends not only on how well the classification functions containing the true class identity perform. If a sample does not receive the full number of votes for the correct class then the outcome may be determined by the voting in the 'redundant' classification functions (i.e. those that do not contain the true class). Although the subsets can be optimised for distinguishing between a given class pair we cannot determine what effect these optimal subsets have when the sample is not from one of these classes. Thus, it is possible that subsets with poorer separability indices can produce better classification results due to their ability to share the 'redundant' votes more evenly amongst the false class identities.

6.9 COMBINING TEXTURE MEASURES

Texture measures from the different analyzers were combined and entered into the variable selection process in an attempt to find feature subsets with greater discriminating power. This was carried out on combined co-occurrence and Laws' mask measures, on combined co-occurrence and

edge pair measures, and finally on the combined texture measures of all three analyzers.

6.9.1 Combined Co-occurrence and Laws' Mask Measures

A combination of eight co-occurrence features and the full eight Laws' mask features were used in the accelerated search procedure. The eight co-occurrence features chosen were the same as those selected from the previous database.

The T^2 statistics indicated that in many cases there was no combination of Laws' mask and co-occurrence features that could exceed the separability managed with their independent use. A stronger two feature subset was discovered in only seven of the fifteen class pairs. However, the classification results displayed in Table 6.15 below do not reflect this fact. The results were consistently better than those obtained with co-occurrence and Laws' mask measures used alone. Obviously the few better subsets that were discovered had a strong influence on the final classification outcome.

Features Used	% Correct Classification
2	77.2
3	89.6
4	89.6
5	91.7
6	90.6
7	90.6
8	89.6

Table 6.15 Classification success using co-occurrence and Laws mask texture measures combined.

6.9.2 Combined Co-occurrence and Edge Pair Measures

These combined measures consisted of the eight co-occurrence features previously selected, and the full twelve edge pair features. The accelerated search routine was used with these data for variable selection.

The T^2 statistics of the selected subsets again revealed that only a few combinations had greater powers of separability. In general the optimal subsets selected for a given discriminant function were composed either of co-occurrence measures, or of edge pair measures. Only three of the fifteen selected subsets combined measures from both analyzers. Furthermore, the T^2 values were in most cases lower than those for the combined Laws' mask and co-occurrence subsets.

The results of classification, shown in Table 6.16, were quite surprising considering the details discussed above. Combining these texture measures was as successful as combining co-occurrence and Laws' mask data. The success rates were considerably better than when the features from each analyzer were used independently.

It appeared that the main reason for this success was the ability of edge pair measures to distinguish between Plantain and the other classes. The edge pair subsets were particularly superior for differentiating between Plantain and Hazel. The co-occurrence measures were selected exclusively for almost all the other discriminant functions. Therefore, although the overall level of separability was not as great as for the combined co-occurrence and Laws' mask measures, the improved results with the Hazel and Plantain samples produced the lower classification error rate.

Features Used	% Correct Classification
2	81.3
3	90.6
4	89.5
5	92.7
6	90.6
7	90.6
8	90.6

Table 6.16 Classification success using co-occurrence and edge pair texture measures combined.

6.9.3 All texture features combined

As a final experiment the measures from all three texture analyzers were combined into a single file. This was entered first into the backward elimination routine to reduce the thirty-five texture descriptors to suboptimal subsets of fifteen features. The accelerated search program then continued with these subsets. This procedure was obviously quite expensive on computer time.

The classification results are displayed in Table 6.17. The success rate with subsets of 2, 3, and 4 variables were lower than previous results. However, with the larger subsets an improvement was produced. The 95% success rate was roughly comparable to that obtained with the previous database using combined co-occurrence and Laws' mask texture measures.

It must be noted that substantially more computation was required to produce these results. This includes the extra preprocessing techniques used, the additional calculations needed to produce the edge pair measures (which were more complex to produce than the statistical

measures), and the extra work involved in the construction of the classifier. Taking all this into consideration there seems to be little advantage in using the structural texture measures, certainly for classifying these particular textures.

Features Used	% Correct Classification
2	72.9
3	84.4
4	86.5
5	93.6
6	94.8
7	94.8
8	91.7

Table 6.17 Classification success using co-occurrence, Laws mask, and edge pair texture measures combined.

CHAPTER SEVEN

SEARCHING FOR POLLEN

7.1 INTRODUCTION

This Chapter describes briefly the efforts that were made to tackle the problem of locating pollen using SEM derived imagery. Unfortunately there was insufficient time to thoroughly investigate this essential aspect of a automated pollen analysis system. The problem remains one to be tackled further in future work. However, the modest progress that was made during this period of study should still prove highly valuable to any person continuing research in this area, and so is reported here primarily for their benefit.

The discussion in Section 4.1.1 has already indicated the need for computer control of the microscope stage, and broadly specified the major problems that need to be tackled. We will be concerned here only with the image processing aspects of a solution.

7.2 SEPARATING OBJECTS FROM THE BACKGROUND

7.2.1 Segmentation by Thresholding

A principal requirement of a working system will be the need to identify the location of all objects in the scene that is viewed. This classic segmentation problem can probably best be tackled by using simple thresholding techniques. SEM imagery typically displays a very dark background on which the objects are defined by their higher tonal levels. We would expect a grey level histogram derived from this type of image to display a bimodal tendency. The lower modal group would represent the dark background pixels, while the higher group would naturally represent the lighter object pixels. Selecting a threshold

grey level between these groups should allow all pixels to be segmented into object and background classes.

Unfortunately this fundamentally simple idea is often complicated by other factors. These have already been discussed in some detail in Sections 2.4.2 to 2.4.5 of this report. Of particular relevance is the fact that simple thresholding of textured regions is often unsatisfactory. The range of intensities associated with the texture lead to 'holes' within an object wherever pixels have values that fall below the threshold level, since these are consequently classified as background points.

A simple solution is provided by Davis et al. (1975). If the image is first processed by a low pass smoothing filter the extreme values within the textured region are suppressed. The underlying principle is that the textured regions have a higher average grey level, when measured over a local area, than the background regions. Producing a locally averaged image enables a single fixed threshold to successfully segment the image.

Unfortunately the size of local averaging necessary for successful segmentation is problematic and depends on the images themselves. Textures with higher levels of local contrast require a larger degree of local averaging. In the tests undertaken it was found that averaging over 11x11 pixel regions was sufficient for all the taxa used in these experiments.

7.2.2 Selecting a Threshold Automatically

In an automated system the selection of a suitable threshold must be undertaken without the need for human assistance. Although the idea of segmentation by thresholding is a simple one, the automatic selection of a satisfactory level is often a non-trivial task (Rosenfeld and Kak, 1976). After considering several approaches the method proposed by Otsu (1979) appeared to hold the greatest promise of success. This was implemented on the PDP 11/23 machine and in conjunction with the 11x11 smoothing filter was found to perform very well. None of the problems described by Kittler and Illingworth (1985) were encountered during the study period. However, future researchers may like to test this method more thoroughly and, if necessary, employ the extra error checking procedures that are proposed by Kittler and Illingworth.

The scheme that is outlined above is capable of producing a binary image where objects are represented as white pixels and the background as black pixels. The sequence of events is illustrated in Plates 7-I to 7-III. An image displaying a single Oak pollen grain is shown in Plate 7-I. The effect of low pass smoothing on this image is shown in Plate 7-II. Finally, the binary image produced by using the automatic thresholding algorithm is displayed in Plate 7-III.

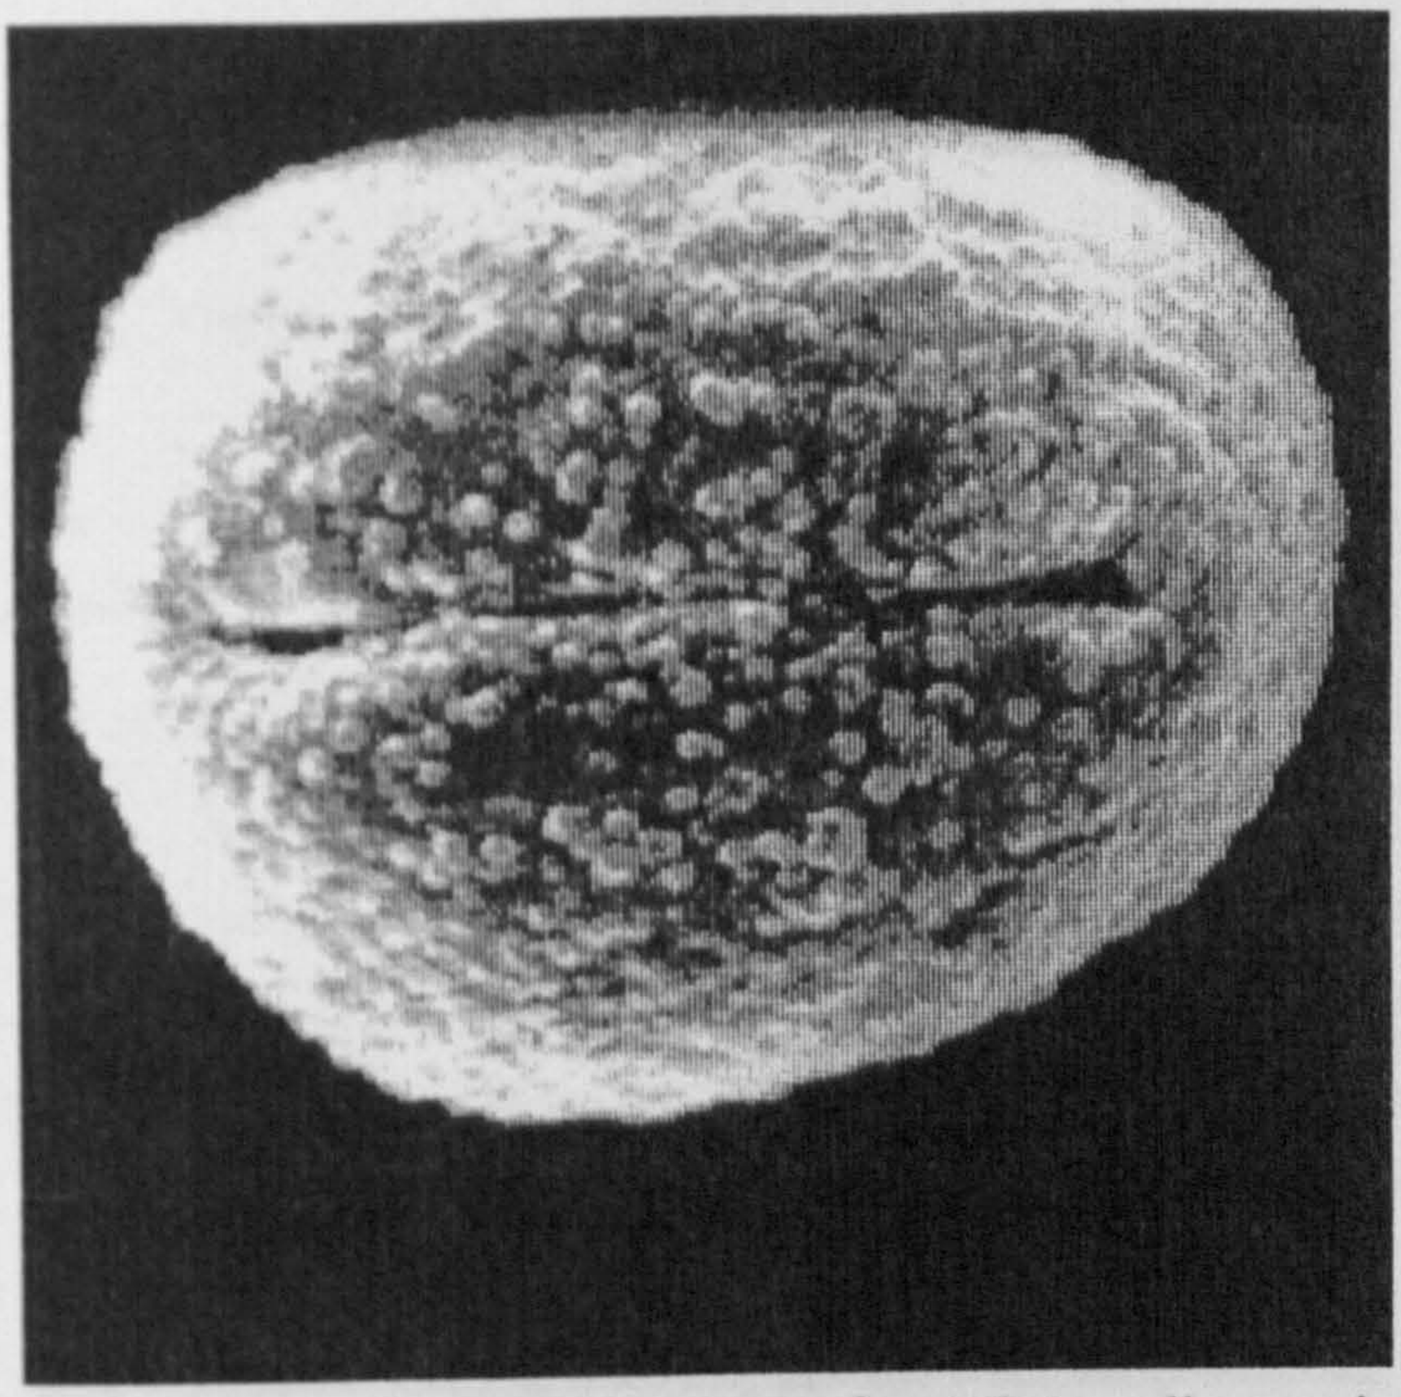

Plate 7-I. Original SEM image of an Oak pollen grain

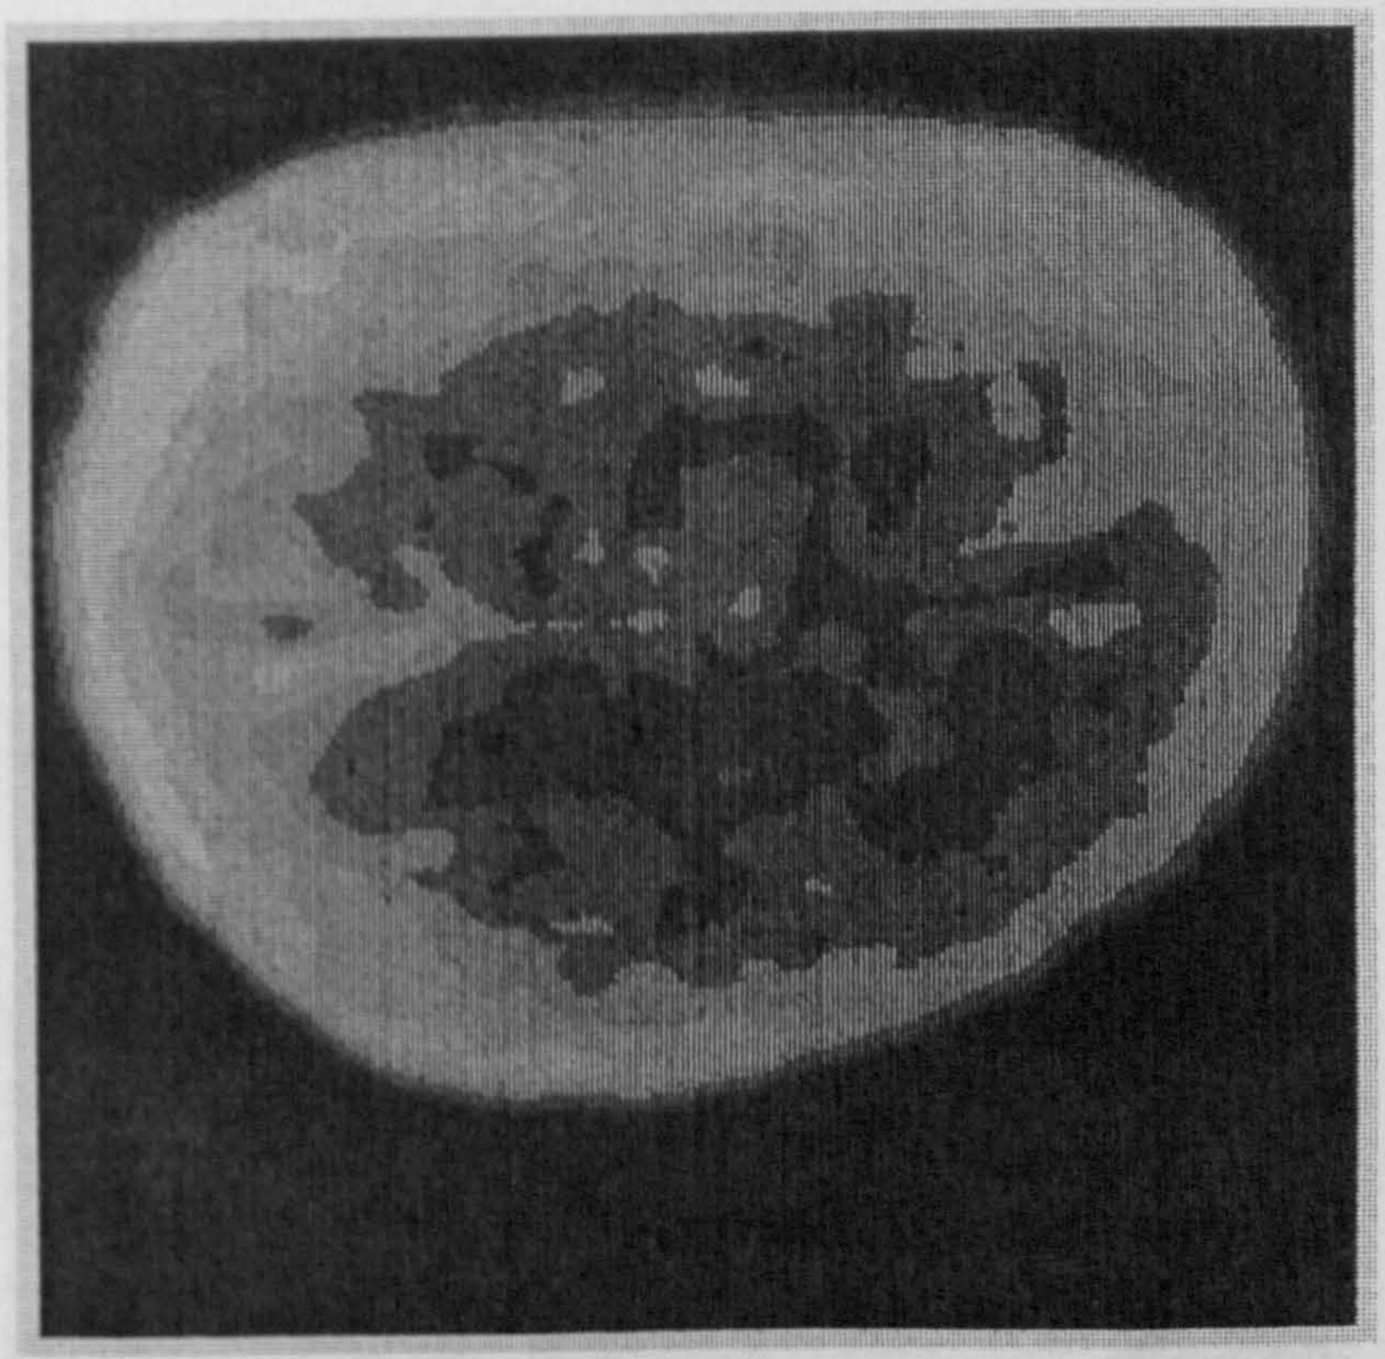

Plate 7-II. Mean smoothed image to suppress texture

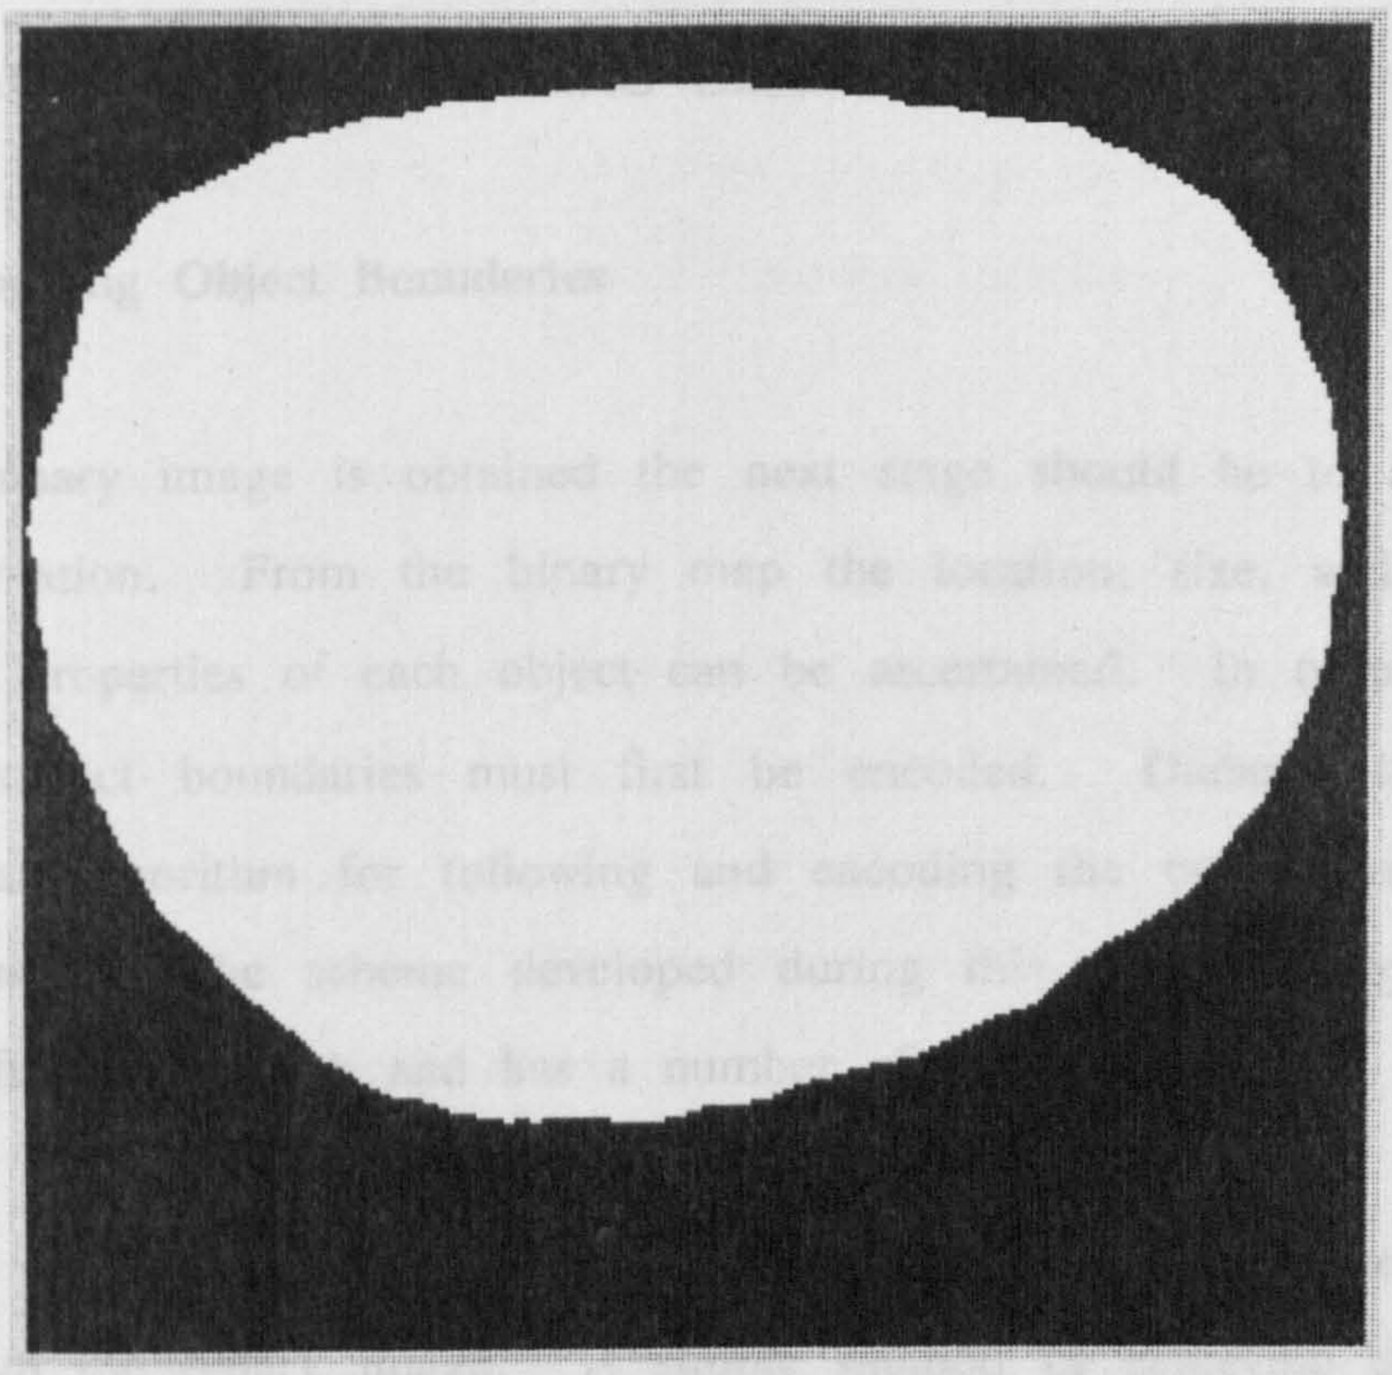

Plate 7-III. Binary image from automatic thresholding

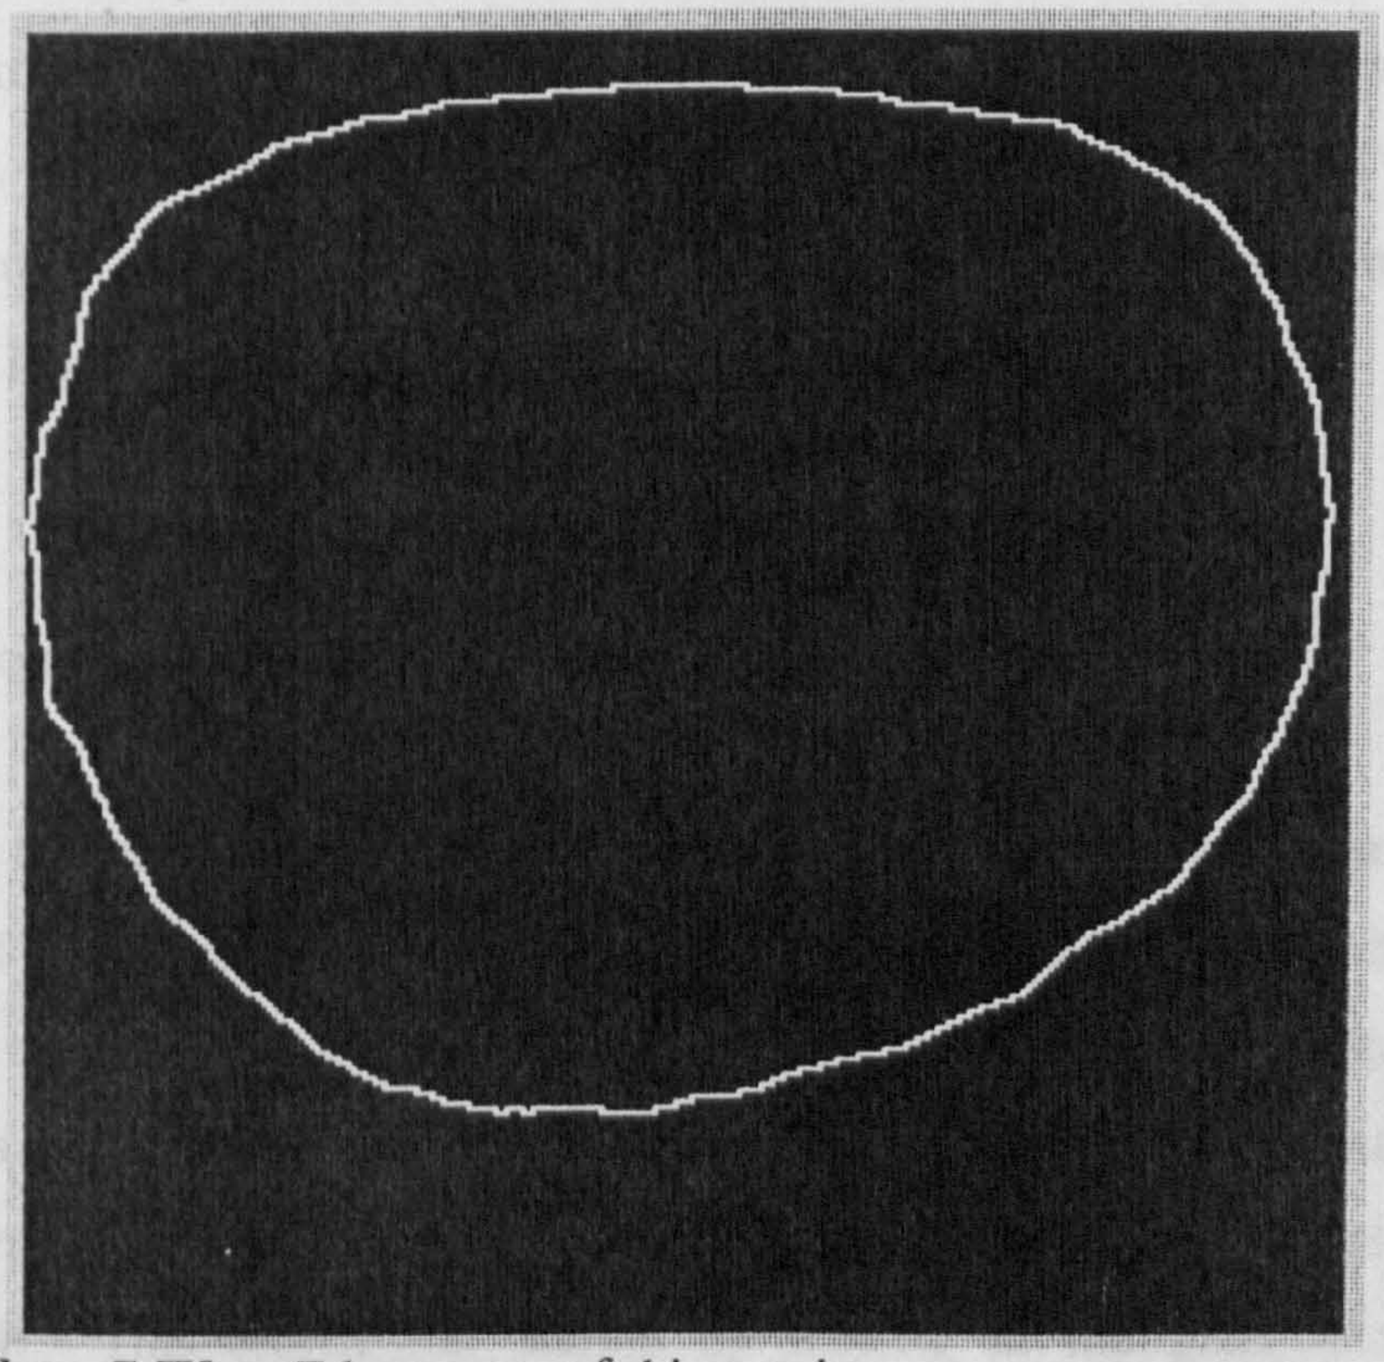

Plate 7-IV. Edge map of binary image

7.3 TRACING AND MEASURING OBJECTS

7.3.1 Locating Object Boundaries

Once a binary image is obtained the next stage should be to analyse this information. From the binary map the location, size, and other geometric properties of each object can be ascertained. In order to do this the object boundaries must first be encoded. Dudani (1976) provides an algorithm for following and encoding the boundaries in a binary image. The scheme developed during this study employed a rather different approach and has a number of advantages.

The first stage in the scheme developed was to derive an edge map image from the binary image. A simple method of achieving this was to process the binary image with the Roberts 2x2 edge operator. Whenever this operator crossed a boundary between black and white regions an edge pixel was generated, but elsewhere zero values were produced. The effect of this action is clearly seen in Plate 7-IV where an edge map of the Oak pollen grain has been produced from the image in Plate 7-III.

7.3.2 Tracing Boundaries

In order to trace object boundaries a search was initiated from the top-left of the edge map image, working left to right along each row, and from one row down to the next. Whenever a nonzero value was found this location became the starting pixel for a boundary following stage. From the initiating pixel a search was made for an adjacent edge pixel. This search was started by looking at the pixel to the left

of the initiating pixel (where an adjacent edge pixel should never be found since it would have become the initiating pixel) and working round all other neighbours in an anti-clockwise direction.

If an edge neighbour was discovered the relative direction of this new edge pixel was recorded using Freeman chain-coding (Freeman, 1961). The search was then continued from the new pixel, starting 45 degrees around from the direction of the old edge pixel. A search was concluded if the boundary ran off the image area, came to a dead-end, or returned to the location of the initiating pixel. All boundaries where chain-encoded as they were traced.

Once a trace had been completed the location of the starting pixel and the chain code of the boundary were used to set all pixels in this trace to zero value. In this way there were no difficulties in avoiding the multiple recording of any given boundary. The search for a new initiating pixel continued from the location of the last initiating pixel, and the process continued in this fashion until the entire image had been searched.

7.3.3 Recording Object Properties

Once the search had been completed the chain codes recorded from object boundaries were analysed. From these data the boundary length and the enclosed area of all completely enclosed objects were determined. These measures were combined to give a 'compactness' index by dividing the square of the perimeter by the enclosed area. The enclosed area calculation provided by Freeman (1961) was developed further to provide the centre of gravity of each object. Therefore, the

location, size, and simple shape characteristics of each object were produced from this procedure.

The corner-finding algorithm of Freeman and Davis (1977) was also briefly experimented with. This was found to produce encouraging results, easily outperforming the methods described by Rosenfeld and Johnston (1973), and Rosenfeld and Wezska (1975), which were also investigated. It appeared to have considerable potential for providing more detailed shape information, but again this remains an area for future research.

7.4 CONCLUSION

This chapter has described the investigations that were undertaken into automatically locating objects in SEM derived imagey, and in identifying pollen from other artifacts. It has described a method to obtain a satisfactory segmentation of objects and background within a given image. The problem of tracing object boundaries has been tackled, using a newly developed algorithm that simplifies the task. Some measures derived from the encoded objects that should be useful in a system to determine where pollen is located have been presented.

Although insufficient experimental results were available to provide a balanced appraisal of the effectiveness of the proposed scheme, the information presented here should prove valuable for future research into this task. Obviously some problems, such as dealing with overlapping grains and non-pollen objects that have a pollen-like shape, have not been addressed but these provide interesting tasks for those undertaking further work.

CHAPTER EIGHT

CONCLUSIONS AND FURTHER WORK

8.1 DIGITAL IMAGE PROCESSING IN PALYNOLOGY

The future for digital image processing within the field of palynology would seem to be secure. Almost all modern scanning electron microscopes capture their image data in digital format, and equally invariably this information is readily exportable into external computer systems. Contrast stretching, thresholding, and frame averaging, are examples of image processing techniques that are now implemented on many SEMs directly. With this easy access to digitally encoded image data the desire for automated computer analysis of the captured information is likely to grow ever stronger.

This study has concentrated on investigating the potential of digital image processing and analysis for the automated recognition of pollen as viewed under a SEM. Before discussing the results that were obtained in greater detail, some statements can be made on a rather broader scale concerning the lessons learnt from this work.

8.2 GENERAL CONCLUSIONS

Digital image analysis offers a method of analysing pollen SEM imagery directly by computer. The potential for automated pollen identification exists provided that suitable software can be developed. Although pollen presents a difficult recognition problem, encouragement may be taken from the fact that successful systems for analysing other biological materials are already in existence.

This study has concentrated on the task of identifying pollen grains, making the assumption that a system for accurately locating grains will

become available in the near future. It has been shown that analyzing the textural properties of the pollen exine is the most likely method of providing measured features that have sufficient reliability and diagnostic power for accurate identification.

The need for sophisticated classification procedures is also clear from the work reported here. It is quite apparent from the results of Chapter Six that classifier design can have a very significant effect on the success of an automated system. Associated with this is the need to select suitable measurements at the feature extraction stage. In order to obtain suitable measurements we require reliable and powerful texture analysis schemes. However, the most suitable features for identifying any given taxa naturally depend on the properties of the taxa itself. Variable selection routines are likely to be vital in establishing the best individual measurements for any given case and so maximize classifier performance.

The attempts described in Chapter Seven at tackling the problem of locating pollen grains within a given scene, although not conclusive, strongly suggest that this will not be an insurmountable difficulty.

8.3 A REVIEW OF THE RESULTS

The success rates attained in Chapter Six are worthy of a more detailed evaluation than is presented above. The results obtained from the first major database of samples clearly illustrated several important points. These were, firstly, that multiple measures of textural properties were essential in order to gain sufficient discriminatory information to allow high levels of classification success. Almost invariably the success rates

increased as more texture measures were utilised. However, it was soon established that an independent assessment of classifier performance is needed in order to prevent biased estimates of the success rate.

With a leave-one-out classification scheme the success rate was found to reach a saturation level as increasing numbers of variables were used. Greater classification accuracy was possible only be resorting to other means, as described below.

Improved results were gained through modifying the original design of the pairwise linear discriminant classifier. The addition of a simple tie-breaker algorithm to resolve samples drawn on votes between two possible classes provided a notable increase in the number of correct identifications. The use of a variable selection procedure combined with a further modification to the classifier in order to capitalize on this extra information, produced even more impressive improvements to success rates.

Using only co-occurrence texture measures in this enhanced classification scheme, success rates in the region of over 90% were produced. Texture features based on local mask matching, the Laws mask texture measures, produced similarly impressive results.

Through studying the results presented in the earlier sections of Chapter Six, it became apparent that the success rates were often 'spoilt' by certain frequently occurring confusions between pollen taxa. Confusions between Rye grass, Elm and Hazel, and between Plantain and Hazel, are typical examples. However, the taxa confused were not consistent between different texture analysis schemes. There was, in effect, confusion in this sense as well.

Combining features from several texture measurement schemes was attempted, using the variable selection routine to guide the choice of combinations. In this way the confusion between certain combinations of taxa ought to be minimized, since features from the most appropriate analysis scheme could be selected.

This approach was quite successful, especially with the second major database of exine samples. Due to changes in the magnification of these samples the statistical texture analysis schemes were less appropriate for quantifying their textural properties. Consequently, the results obtained using the analysis schemes individually show a marked deterioration. However, by combining texture features from several analysis schemes, including those from a structural texture analyzer better suited to these sample, the success rate recovered to over 90% again.

Overall, it is clear that substantial improvements to unbiased estimates of classification success were brought about through improved classifier design, variable selection, and combining texture measures from several analysis schemes.

8.4 THE FUTURE FOR AUTOMATED POLLEN ANALYSIS

The classification success rates attained in Chapter Six of this report suggest that automatic pollen identification using digital image processing is a feasible proposition. There should be no illusion that the production of a satisfactory working system will be a simple task; many difficulties remain that have yet to be tackled. However, we may cautiously anticipate the day when fully automated pollen analysis

becomes a reality.

8.5 AREAS FOR FUTURE WORK

There are a number of readily identified areas that future work should address. The most obvious of these is the need for a system that can accurately locate pollen within the captured scenes.

Computer control of a microscope stage is a problem to be solved through hardware rather than software. It is unlikely that a satisfactory system for this task could not easily be manufactured given sufficient funds. Locating objects in general within captured scenes should not prove to be an especially difficult task. A scheme based on simple automated thresholding, like that proposed in Chapter Seven of this report, should be sufficient to segment images into background and object points.

Differentiating between genuine pollen and other unwanted artifacts may present a tougher problem. Use might be made of the 'elemental analyzer' found on modern SEM equipment. This can provide a break down of the chemical elements in a targeted object. With this information a quick and reliable method to distinguish between organic and non-organic material ought to be possible. Object shape and surface texture descriptions could supplement this information if necessary. Careful preparation of the original samples ought to ensure that very little other organic materials are present.

Preparation techniques are likely to play a valuable role in the production of a satisfactory system. If methods of producing very clean

samples are employed, and if the distribution of pollen can be tightly controlled such that individual grains are well separated then the overall task will be greatly simplified. Therefore, this is another area of interest for future research.

Consideration will need to be given to improving the techniques established in this work. In particular the classification structure could be investigated further. If it is necessary to identify many more pollen taxa than attempted here the pairwise classifier design will be an inappropriate instrument for the task. Possible solutions are the use of multiple discriminant functions, multinomial logit models, or hierarchical based classification systems.

Many other techniques for the quantification of texture exist, and some may prove more effective for pollen identification than those investigated here. A feature selection routine, like that described in Chapter Five, is likely to be useful in determining the power of new texture measures, and would also be invaluable when constructing a hierarchical classifier.

REFERENCES

Abdou, I.E. and Pratt, W.K. (1979) "Quantitative design and evaluation of enhancement/thresholding edge detectors", *Proceedings of the IEEE*, vol. 67, 753-763

Adams, R.J. and Morton, J.K. (1972) "An improved technique for examining pollen under the scanning electron microscope", *Pollen Et Spores*, vol. 14, no. 2, 203-212

Ahuja, N. and Rosenfeld, A. (1978) "A note on the use of second-order gray-level statistics for threshold selection", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 12, 895-898

Alparslan, E. and Ince, F. (1981) "Image enchancement by local histogram stretching", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-11, no. 5, 376-385

Arazi, B. (1977) "Handwriting identification by means of run-length measurements", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-7, no. 12, 878-881

Ataman, E., Aatre, V.K. and Wong, K.M. (1980) "A fast method for real time median filtering", *IEEE Transactions on Acoustics, Speach and Signal Processing*, vol. ASSP-28, 415-420

Bailey, D.G. and Hodgson, R.M. (1985) "Range filters: Local-intensity subrange filters and their properties", *Image and Vision Computing*, vol. 3, no. 3, 99-110

Batchelor, B.G., and Marlow, B.K. (1980) "Fast generation of chain code", *IEE Proceedings*, vol. 127, part E, no. 4, 143-147

Bednar, J.B. and Watt, T.L. (1984) "Alpha trimmed means and their relationship to median filters", *IEEE Transactions on Acoustics, Speach and Signal Processing*, vol. ASSP-32, 145-153

Bennet, J.R. and Mac Donald, J.S. (1975) "On the measurement of curvature in a quantized environment", *IEEE Transactions on Computers*, vol. C-24, no. 8, 803-820

Bernstein, R. (Ed.) (1978) *Digital Image Processing for Remote Sensing.* New York: IEEE Press, 473p

Bertolini, L. and Vernazza, G. (1982) "Electromagnetic effect evaluation by Markovian texture analysis of nucleated cells", in Levialdi, S. (ed.), *Digital Image Analysis.* London: Pitman, 388p

Brasier, M.D. (1980) *Microfossils.* London: Allen and Unwin, 193p

Caratini, C. (1981) "Ultrasonic sieving using an efficient device to improve substantially our palynological activities", *Pre-conference abstracts of the 5th International Conference of Palynology*, Cambridge, 67

Case, S.K., Almeida, S.P., Dallas, W.J., Fournier, J.M., Pritz, K., Cairns, J., Jr., Dickson, K.L., and Pryfogle, P.A. (1978) "Coherent microscopy and matched spatial filtering for real-time recognition of diatom species", *Environmental Science and Technology*, vol. 12, no. 8, 940-946

Castleman, K.R (1979) *Digital Image Processing.* Englewood Cliffs (N.J.): Prentice-Hall, 429p

Chen, P.C. and Pavlidis, T. (1979) "Segmentation by texture using a

co-occurrence matrix and a split-and-merge algorithm", *Computer Graphics and Image Processing*, vol. 10, 172-182

Chien, Y.P. and Fu, K.S. (1974) "Recognition of X-ray picture patterns", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-4, no. 3, 145-156

Curran, P.J. (1985) *Principles of Remote Sensing.* London: Longman, 282p

Cusma Velari, T. (1984) "Effects of preparation techniques on pollen grains of *Centaurea weldeniana* (Asteraceae)", *Grana*, vol. 23, 91-95

Damblon, F. (1975) "Sputtering, a new method for coating pollen grains in scanning electron microscopy", *Grana*, vol. 15, 137-144

Davis, L.S. (1977) "Understanding shape: angles and sides", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-26, no. 3, 236-242

Davis, L.S., Clearman, M., and Aggarwal, J.K. (1981) "An empirical evaluation of generalized co-occurrence matrices", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-3, no. 2, 214-221

Davis, L.S., Johns, S.A., and Aggarwal, J.K. (1979) "Texture analysis using generalised co-occurrence matrices", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-1, no. 3, 251-259

Davis, L.S., Rosenfeld, A. and Weszka, J.S. (1975) "Region extraction by averaging and thresholding", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-5, 383-388

Dickson, K.L., Slocomb, J.P., Cairns, J., Jr., Almeida, S.P., Eu, J.K.T. (1977) "A laser-based optical filtering system to analyze samples of diatom communities", in Cairns, J., Jr., Dickson, K.L., and Westlake, G.F., (eds.), *Biological Monitoring of Water and Effluent Quality.* American Society for Testing and Materials, 121-132

Dinstein, I. and Shapira, Y. (1982) "Ancient Hebraic handwriting identification with run-length histograms", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-12, no. 3, 405-409

Don, H., Fu, K., Liu, C.R. and Lin, W. (1984) "Metal surface inspection using image processing techniques", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-14, no. 1, 139-146

Dudani, S.A. (1976) "Region extraction using boundary following", in Chen, C.H. (ed.), *Pattern Recognition and Artificial Intelligence.* New York: Academic Press, 216-232

Dyer, C.R., Hong, T. and Rosenfeld, A. (1980) "Texture classification using gray level co-occurrence based on edge maxima", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-10, no. 3, 158-163

Faegri, K. and Iversen, J. (1975) *Textbook of Pollen Analysis.* 3rd Edition by Faegri, K., Oxford: Blackwell, 295p

Fram, J.R. and Deutsch, E.S. (1975) "On the quantitative evaluation of edge detection schemes and their comparison with human performance", *IEEE Transactions on Computers*, vol. C-24, no. 6, 616-628

Freeman, H. (1961) "On the encoding of arbitrary geometric configurations", *IRE Transactions on Electronic Computers*, vol. EC-10, 260-268

Freeman, H. and Davis, L.S. (1977) "A corner-finding algorithm for chain-coded curves", *IEEE Transactions on Computers*, vol. C-26, no. 3, 297-303

Forster, R.M. (1986) "Palynomorph purification and fractionation by equilibrium density gradient centrifugation", University of Hull, Theses: Geography. England. 146p

Galloway, M. (1974) "Texture analysis using gray level run lengths", *Computer Graphics and Image Processing*, vol. 4, 172-199

Germeraad, J.H. and Muller, J. (1970) "A computer based numerical coding system for the description of pollen grains and spores", *Review of Palaeobotany and Palynology*, vol. 10, 175-202

Germeraad, J.H., and Muller, J. (1971) *A Computer Based Numerical Coding System for the Description of Pollen Grains and Spores.* Leidon: Rijksmus Geol. Mieral, vol. I-II

Gerson, D.J. and Rosenfeld, A. (1975) "Automatic sea ice detection in satellite pictures", *Remote Sensing of Environment*, vol. 4, 187-198

Gonzalez, R.C. and Wintz, P. (1977) *Digital Image Processing.* London: Addison-Wesley, 431p

Gordon, A.D. (1981) *Classification.* London: Chapman and Hall, 193p

Guppy, J., Milne, P., Glikson, M. and Moore, M. (1973) "Further developments in computer assistance to pollen identification", *Special Publication no. 4, Geological Society of Australia*, 201-206

Hall, E.L. (1979) *Computer Image Processing and Recognition.* New York: Academic Press, 584p

Hand, D.J. (1981) *Discrimination and Classification.* Chichester: Wiley, 218p

Haralick, R.M. (1974) "A measure for circularity of digital figures", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-4, no. 7, 394-396

Haralick, R.M. (1979) "Statistical and structural approaches to texture", *Proceedings of the IEEE*, vol. 67, no. 5, 786-804

Haralick, R.M. and Shanmugam, K. (1973) "Computer classification of resevoir sandstones", *IEEE Transactions on Geoscience and Electronics*, vol. GE-11, 171-177

Haralick, R.M., Shanmugam, K. and Dinstein, I. (1973) "Textural features for image classification", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-3, no. 3, 610-621

Heusser, L.E. and Stock, C.E. (1984) "Preparation techniques for concentrating pollen from marine sediments and other sediments with low pollen density", *Palynology*, vol. 8, 225-227

Hill, M.O. (1979a) "TWINSPAN - A FORTRAN program for arranging multivariate data in a ordered two-way table by classification of the individual and attributes", *Ecology and Systematics*, Cornell University, New York

Hill, M.O. (1979b) "DECORANA - A FORTRAN program for detrended correspondance analysis and reciprocal averaging", *Ecology and Systematics*, Cornell University, New York

Hodgson, R.M., Bailey, D.G., Naylor, M.J., Ng, A.L.M., and McNeill, S.J. (1985) "Properties, implementations and applications of rank filters", *Image and Vision Computing*, vol. 3, no. 1, 3-14

Hong, T., Dyer, C.R., and Rosenfeld, A. (1980) "Texture primitive extraction using an edge-based approach", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-10, no. 10, 659-675

Hong, T., Wu, A.Y., and Rosenfeld, A. (1980) "Feature value smoothing as an aid in texture analysis", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-10, no. 8, 519-524

Huang, T.S., Yang, G.J., and Tang, G.Y. (1979) "A fast two-dimensional median filtering algorithm", *IEEE Transactions on Acoustics, Speach, and Signal Processing*, vol. ASSP-27, no. 1, 13-18

Irons, J.R., and Petersen, G.W. (1981) "Texture transforms of remote sensing data", *Remote Sensing of Environment*, vol. 11, 359-370

Jarvis, J.F. and Roberts, C.S. (1976) "A new technique for displaying continuous tone images on a bilevel display", *IEEE Transactions on Communications*, vol. COM-24, 891-898

Julez, B. (1962) "Visual pattern discrimination", *IRE Transactions on Information Theory*, vol. 8, no. 2, 84-92

Kakikura, M. (1979) "A simple method of extracting the feature axes of objects", *Bulletin Electrotechnical Laboratory (Japan)*, vol. 43, no. 6, 351-357

Kanal, L. (1974) "Patterns in pattern recognition: 1968-1974", *IEEE Transactions on Information Theory*, vol. IT-20, no. 6, 697-722

Kim, C.E. and Strintzis, M.G. (1980) "High speed multidimensional convolution", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-2, 269-273

Kirby, R.L. and Rosenfeld, A. (1979) "A note on the use of (gray level, local average gray level) space as an aid in threshold selection", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-9, no. 12, 860-864

Kittler, J. and Illingworth, J. (1985) "On threshold selection using clustering criteria", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-15, no. 5, 652-655

Kruger, R.P., Tompson, W.B. and Turner, A.F. (1974) "Computer diagnosis of pneumoconiosis", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-4, no. 1, 40-49

Langford, M., Taylor, G.E., and Flenley J.R. (1986) "The application of texture analysis for automated pollen identification", *Conference on Identification and Pattern Recognition*, vol. II, RAI/IPAR'86, Toulouse, France, 729-739

Lee, C.C. (1983) "Elimination of redundant operations for a fast Sobel operator", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, 242-245

Leffingwell, H.A. and Hodkin, N. (1971) "Techniques for preparing fossil palynomorphs for study with the scanning and transmission electron microscopes", *Review of Palaeobotany and Palynology*, vol. 11, 171-199

Lillesand, T.M. and Kiefer, R.W. (1987) *Remote Sensing and*

Image Interpretation. 2nd Edition, London: Wiley, 721p

Lynch, S.P. and Webster, G.L. (1975) "A new technique of preparing pollen for scanning electron microscopy", Grana, vol. 15, 127-136

Mannion, A.M. (1980) *Pollen Analysis.* Reading: University of Reading, 62p

Mason, D.C. (1979) "Segmentation of terrain images using textural and spectral characteristics", *Computers and Digital Techniques*, vol. 2, no. 6, 251-259

McFarlane, M.D. (1972) "Digital pictures fifty years ago", *Proceedings of the IEEE*, vol. 60, no. 7

Mirkin, G.R. and Bagdasaryan, L.L. (1972) "The feasibility of identifying paleontological objects with the aid of optical analyzing systems", *Paleontological Journal*, vol. 6, 103-108

Mitchell, O.R., Myers, C.R., and Boyne, W. (1977) "A maximum measure for image analysis", *IEEE Transactions on Computers*, vol. C-25, 408-414

Moore, P.D. and Webb, J.A. (1983) *An Illustrated Guide to Pollen Analysis.* London: Hodder and Stoughton, 133p

Morrin, T.H. (1974) "A black-white representation of a gray scale picture", *IEEE Transactions on Computers*, vol. C-23, 184-186

Nagao, M. and Matsuyama, T. (1979) "Edge preserving smoothing", *Computer Graphics and Image Processing*, vol. 9, 394-407

Nair, P.K.K. (1966) *Essentials of Palynology.* London: Asia Publishing House, 96p

Nakagawa, Y. and Rosenfeld, A. (1978) "A note on the use of local Min and Max operations in digital picture processing", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 8, 632-635

Niblack, W. (1986) *An Introduction to Digital Image Processing.* London: Prentice-Hall, 215p

Nilsson, S., Nybom, R. and Praglowski, J. (1974) "Experiments regarding collapsing of pollen grains in scanning electron microscopy", *Grana*, vol. 14, 23-25

O'Handley, D.A. and Green, W.B. (1972) "Recent developments in digital image processing at the image processing laboratory at the Jet Propulsion Laboratories", *Proceedings of the IEEE*, vol. 60, 821-828

Otsu, N. (1979) "A threshold selection method from gray-level histograms", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-9, no. 1, 62-66

Pankhurst, R.J. (1978) *Biological Identification.* London: Edward Arnold, 104p

Peleg, S. (1978) "Iterative histogram modification 2", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 7, 555-556

Pietikainen, M.K., and Rosenfeld, A. (1982) "Edge-based texture measures", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-12, no. 4, 585-594

Pietikainen, M.K., Rosenfeld, A. and Davis, L.S. (1983) "Experiments with

texture classification using averages of local pattern matches", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 3, 421-426

Ridler, T.W. and Calvard, S. (1978) "Picture thresholding using an iterative selection method", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 8, 630-632

Roetling, P.G. (1976) "Halftone method with edge enhancement and Moire suppression", *Journal of the Optical Society of America*, vol. 66, no. 10, 985-989

Roetling, P.G. (1977) "Binary approximations to continuous tone images", *Photographic Science and Engineering*, vol. 21, no. 2, 60-65

Rosenfeld, A. (1974a) "Compact figures in digital pictures", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-4, no. 4, 221-223

Rosenfeld, A. (1974b) "Digital straight line segments", *IEEE Transactions on Computers*, vol. C-23, no. 12, 1264-1269

Rosenfeld, A. (1975) "A note on automatic detection of texture gradients", *IEEE Transactions on Computers*, vol. C-24, no. 10, 988-991

Rosenfeld, A. and Davis, L.S. (1978) "Iterative histogram modification", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 4, 300-302

Rosenfeld, A. and Davis, L.S. (1979) "Image segmentation and image models", *Proceedings of the IEEE*, vol. 67, no. 5, 764-772

Rosenfeld, A. and Johnston, E. (1973) "Angle detection on digital curves", *IEEE Transactions on Computers*, vol. C-22, 875-878

Rosenfeld, A. and Kak, A.C. (1976) *Digital Picture Processing.* 1st Edition, New York: Academic Press, 457p

Rosenfeld, A. and Thurston, M. (1972) "Edge and curve detection for digital scene analysis", *IEEE Transactions on Computers*, vol. C-20, no. 5, 562-569

Rosenfeld, A., Wang, C., and Wu, A.Y. (1982) "Multispectral texture", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-12, no. 1, 79-84

Rosenfeld, A. and Wezska, J.S. (1975) "An improved method of angle detection on digital curves", *IEEE Transactions on Computers*, vol. C-24, no. 9, 940:941

Rowley, J.R. (1973) "Formation of pollen exine bacules and microchannels on a Glycocalyx", *Grana*, vol. 13, 129-138

Rutovitz, R., Green, D.K., Farrow, A.S.J., and Mason, D.C. (1978) "Computer-assisted measurement in the cytogenetic laboratory", p301-329, in Batchelor, B.G. (ed.), *Pattern Recognition: ideas in practice*, New York: Plenum Press, 485p

Sadjadi, F.A. and Hall, E.L. (1978) "Numerical computation of moment invarients for scene analysis", *Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing.*

Schowengerdt, R.A. (1983) *Techniques for Image Processing and Classification in Remote Sensing.* London: Academic Press, 249p

Shen, H.C., and Wong, A.K.C. (1980) "Generalized texture representation and

metric", IEEE, 695-703

Sklansky, J. (1978) "Image segmentation and feature extraction", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 4, 237-247

Squires, R.H. (1970) *A Computer Program for the Presentation of Pollen Data.* Durham: University of Durham, 53p

Starkey, J. and Simigian, S. (1987) "IMAGE: A FORTRAN V program for image analysis of particles", *Computers and Geosciences*, vol. 13, no. 1, 37-59

Steinberg, D.I. (1974) *Computational Matrix Algebra.* New York: McGraw-Hill, 280p

Stewart, G.W. (1973) *An Introduction to Matrix Computations.* New York: Academic Press, 441p

Stoffel, J.C. (1982) *Graphical and Binary Image Processing and Applications.* Mass: Artech House, 580p

Stoffel, J.C. and Moreland, J.F. (1981) "A survey of electronic techniques for pictorial reproduction", *IEEE Transactions on Communications*, vol. COM-29, no. 12, 1898-1925

Sutton, R.N. and Hall, E.L. (1972) "Texture measures for automatic classification of pulmonary disease", *IEEE Transactions on Computers*, vol. C-21, 667-676

Tamura, H., Mori, S. and Yamawaki, T. (1978) "Textural features corresponding to visual perception", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 6, 460-473

Thomas, I.L., Benning, V.M. and Ching, N.P. (1987) *Classification of Remotely Sensed Images.* Bristol: Adam Hilger, 268p

Ting, D. and Prasada, B. (1974) "Digital processing techniques for encoding of graphics", *Proceedings of the IEEE*, vol. 6, 751-769

Tomita, F. and Tsuji, S. (1977) "Extraction of multiple regions by smoothing in selected neighborhoods", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-7, 107-109

Tomlinson, P. (1984) "Ultrasonic filtration as an aid in pollen analysis of archaeological deposits", *Circaea*, vol. 2, no. 3, 139-140

Trussel, H.J. (1979) "Comments on 'Picture thresholding using an iterative selection method'", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-9, no. 5, 311

Umetani, Y. and Taguchi, K. (1979) "Feature properties to discriminate complex shapes", *9th International Symposium on Industrial Robots*, 367-378

Van der Hammen, T., Werner, J.H., and Van Dommolen, H. (1973) "Palynological record of the upheaval of the northern Andes; A study of the Pliocene and Lower Quaternary of the Columbian Eastern Cordillera and the early evolution of its high-Andean biota", *Review of Palaeobotany and Palynology*, vol. 16, 1-122

Walker, D., Milne, P., Guppy, J. and Williams, J. (1968) "The computer assisted storage and retrieval of pollen morphological data", *Pollen Et Spores*, vol. 10, 251-262

Walker, P.A. and Grant, I.W. (1986) "QUADTREE: A FORTRAN program to extract the quadtree structure of a raster format multicolored image", *Computers and Geosciences*, vol. 12, no. 4A, 401-410

Wang, S., Dias Velasco, F.R., Wu, A.Y. and Rosenfeld, A. (1981) "Relative effectiveness of selected texture primitive statistics for texture discrimination", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-11. no. 5, 360-370

West, R.G. (1977) *Pleistocene Geology and Biology.* 2nd Edition, London: Longmans, 440p

Weszka, J.S. (1978) "A survey of threshold selection techniques", *Computer Graphics and Image Processing*, vol. 8, 259-265

Weszka, J.S., Dyer, C.R. and Rosenfeld, A. (1976) "A comparative study of texture measures for terrain classification", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-6, no. 4, 269-285

Weszka, J.S., Nagel, R.N. and Rosenfeld. A. (1974) "A threshold selection technique", *IEEE Transactions on Computers*, vol. C-23, 1322-1326

Weszka, J.S. and Rosenfeld, A. (1978) "Threshold evaluation techniques", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-8, no. 8, 622-629

Weszka, J.S. and Rosenfeld, A. (1979) "Histogram modification for threshold selection", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-9, no. 1, 38-52

Wong, A.K.C. and Vogel, M.A. (1977) "Resolution dependent information measures for image analysis", *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-7, no. 1, 49-61

Wong, R.Y. and Hall, E.L. (1978) "Scene matching with invariant moments", *Computer Graphics and Image Processing*, vol. 8, 16-24