2409

Ancient Monuments Laboratory
Report 27/94

AN ASSESSMENT OF DIGITAL IMAGE PROCESSING IN CONSERVATION

Dr Chris Caple and Phil Clogg

Ancient Monuments Laboratory Report 27/94

AN ASSESSMENT OF DIGITAL IMAGE PROCESSING IN CONSERVATION

Dr Chris Caple and Phil Clogg

Summary

The emergence of powerful, low cost computer hardware has made Digital Image Processing (DIP) more widely available to institutions and individuals. A detailed assessment was undertaken to determine the potential of these techniques within Conservation particularly for the enhancement of X-radiographs. The equipment and theoretical aspects of digital image processing are reviewed and a series of practical guidelines are described which form the basis for a structured approach to digital image processing of X-radiographs.

Authors' address :-

Dr Chris Caple and Phil Clogg

Department of Archaeology
University of Durham
46 Saddler Street
Durham
DH1 3NU

TABLE OF CONTENTS

LIST OF FIGURES

1. INTRODUCTION

X-radiography has become a technique used extensively throughout conservation for the non destructive examination of artefacts and materials. The ability of the technique to reveal 'non visible' features such as surfaces, constructional details, compositional differences etc. provides a great potential for the study of the form, condition and construction of artefacts and materials.

In order to recover the maximum amount of information the highest quality radiograph is necessary. This can be defined as one which most clearly shows thickness and density changes within the artefact and has a wide contrast range over the area of interest. A range of well established techniques and mechanisms, documented in standard radiographic textbooks, are available to control radiographic quality (Eastman Kodak 1969) and aspects of radiography and interpretation are covered within existing conservation training courses.

The representation of a three dimensional object complete with internal details, superimposed surfaces and depth (characterised by the degree in image sharpness) in two dimensions can produce a very complex image. Thus even the highest quality radiograph can cause difficulties in interpretation. Further processing is therefore often desirable in order to clarify features and details within the image.

Photographic techniques may be used in order to adjust the contrast or magnify elements within the image and a number of more specialised techniques for improving radiographic imaging have been developed within Medicine and Industry. Micro focal systems are available for producing magnified images directly onto the recording medium and are the basis for the developing 'real time' X-ray imaging. Computer Aided Tomography (CAT) is a well established technique within medicine for generating three-dimensional images or for selective imaging of features within a mass. Xeroradiography is a variation of X-radiography producing a different form of image with enhanced edge detail.

1.0.0.1. The present limitations of use

At present conservation radiography is generally limited to the use of standard techniques for the production of images with occasional recourse to photographic enhancement or microscopic enlargement. This is considered adequate for virtually all applications and whilst the majority of users would accept the potential of radiography it is usually perceived as merely an aid to visual identification rather than as an analytical tool in its own right. Little has therefore been published on the interpretation of conservation radiographs (Corfield 1982) and no textbook on radiographs of archaeological materials has yet been produced.

The infrequent use of the more advanced techniques such as tomography and micro focal radiography is due not only to the high cost and restricted availability of the units but also to the lack of awareness of what additional information may be gained from using these processes. This may be viewed as a function of the lack of research in conservation as no-one, with the exception of Leo Biek, has consistently tried to explore what additional information may be gained from archaeological material using radiographs (Biek 1969).

1.0.0.2. Potential developments

The use of radiography purely as an aid to the visual identification of an artefact combined with the apparent lack of exploitation of the technique within conservation raises concern over the amount of radiographic information which is routinely being overlooked.

Radiography has the advantage of being a relatively rapid, non destructive technique which is used routinely throughout archaeological conservation. The film generally in use has two layers of emulsion and can have a grain size as small as 0.25 microns (Agfa D2). It is therefore capable of recording the finest detail over a wide range of image densities making the radiographic image a highly accurate record of the variations in structure (internal and external) of an object and of any associated material within its immediate surroundings (e.g. outer corrosion layers and adhering mineral material). Difficulties in the interpretation of radiographs arise due to the representation of this three dimensional data in two dimensions and the ability of the radiographic film to hold more data than is perceived by the human visual system.

There is clearly a need for a process by which radiographs can be enhanced to facilitate a more positive interpretation of features in addition to being scanned for data which may not be immediately discernible by the conservator. Photographic techniques have some limited application in this area however the process can be time consuming and relies on the skill of the operator (who may not be the conservator) to extract the information. Advanced techniques such as computer aided tomography and micro focal systems present some means of image manipulation however the previously mentioned disadvantages (high cost, restricted availability etc.) rule out these techniques for routine work. Real time X-ray systems provides for rapid, detailed examination of materials with the results often being recorded in directly in digital format. High resolution output however must still rely on radiographic film and the high cost of the units is a major restriction.

The advent of computers and digital technology has realised the recording and storage of images in digital form thereby making possible an interactive approach to image manipulation. This approach was pioneered with the space programmes in the 1960's to improve images from the space probes and is now used routinely in a number of areas such as medical radiography - e.g. computer aided tomography (CAT), and in geography (GIS) for the study of land use and pollution problems from satellite imagery.

The recent advances in digital technology has made computers generally more accessible thereby opening up digital technology to new areas and applications. Initial work concerned with the acquisition of digital images from standard radiographic film and their subsequent processing has been undertaken by undergraduates at both the Universities of Cardiff and Durham.

In 1991 Nico Morgan carried out research in Durham to explore the feasibility of digitising images from conventional radiographs (Morgan 1991). The work was undertaken on an imaging system based in the Department of Engineering and produced a number of enhanced images the most successful of which featured the addition of false colour to a radiograph effecting a three dimensional representation of the object (see figure 30).

In Cardiff Eric Lange worked with Visual Logic Systems, a commercial software company with some expertise in enhancing medical radiographs, to examine the use of digital image processing on archaeological radiographs. This work produced some positive conclusions (Lange & Watkinson 1992) and resulted Visual Logic Systems extending their service to include archaeological radiographs.

A subsequent piece of work undertaken by Henry Johnson in Durham (Johnson 1992) explored the possibility of calculating various object parameters such as thickness and type of material by using the intensity of measured radiation. The work was carried out in part using a real time X-ray facility (X-Tec Ltd.) to record the images.

These initial pieces of work revealed the ease with which standard radiographic images could be digitised and suggested possibilities for their subsequently manipulation and enhancement. This report aims to provide a comprehensive introductory evaluation of digital image processing as applied to archaeological radiographs and its potential use within archaeological conservation.

The work has been approached with the confidence that advances in computer technology and digital information will provide wider access to high resolution images, at a lower cost, in the near future. It is therefore concerned purely with the evaluation of the techniques and the potential for digital image processing in furthering the examination of archaeological artefacts.

Digital image processing is concerned with the manipulation of digitised images in order to extract more detailed information for human interpretation or for data analysis. The techniques used can be divided into two separate areas depending upon the type of data required. One area is concerned with the enhancement of images in order to provide clearer information on the scene or composition of the image whilst the second area is concerned with the quantification and appraisal of features within the image i.e. the analysis of the image objects. Whilst the two areas and techniques are not entirely discrete it is a convenient division when considering digital image processing. This report concerns itself only with the area of image enhancement and should be considered as purely an introduction to the techniques available and the principles on which image enhancement is based. A guide or protocol for obtaining good quality digitised radiographic images is presented and a structured approach to the choice and application of techniques is suggested (section four).

In addition to the illustrations within the text a selection of images are provided on disk enabling more detailed viewing and manipulation. Appendix 2 describes and lists a number of look up tables (LUT) which have been used within the illustrations and these are provided on disk as text files in a format suitable for the Foster Findlay PC_Image Plus program. Appendix 3 describes a series of program modules designed to explore some specifically designed algorithms which have proved particularly relevant for the enhancement of radiographs. These are provided on disk and run under Windows with Foster Findlay PC_Image Plus. Finally a short demonstration program is included illustrating some image processing techniques, details of which are given in appendix 5. It is hoped that this report and its supporting material will provide a sound base for those working in conservation to further explore the specific requirements of image enhancement of radiographs.

All the discussion of image processing techniques in this report are concerned only with the processing of monochrome and pseudo or false colour images. Full colour image processing was not considered appropriate for this pilot study due to the grey scale nature of the radiographic image as well as the cost. Similarly the choice of image resolution was based on considerations of processing time, storage space and cost.

It should be noted that any discussion of computers and ancillary equipment is in relation to the IBM PC compatible type machines and operating systems. Many other systems are available (MacIntosh, UNIX etc.) and these will have there own specific requirements in the areas of hardware and software. The processing techniques however remain the same which ever system is chosen.

The nature of a computer based technique necessitates some discussion of the type of equipment required and a minimum powered image processing set-up is suggested (Appendix one).

2. A DIGITAL IMAGING SYSTEM

2.1. THE DIGITAL IMAGE

A digital image consists of a two dimensional array of numbers representing the light intensity at each point in a set of two dimensional spatial co-ordinates. Each point within the array is called a pixel (or pel) being derived from the term 'picture-element'. Each pixel defines the level of brightness (the grey level) at a particular co-ordinate. Pixels are normally rectangular in shape - the ratio of the sides of the rectangle dictating the aspect ratio of the pixel. The size of the two dimensional array of pixels is termed the spatial resolution of the image. This would typically be at least 640 x 480 pixels which approximates to the resolution of a television image. The number of grey levels available is termed the brightness resolution and requires a scale ranging from black to white. It is apparent that 2^8 or 256 discrete grey levels are sufficient to represent well the information present in most images. This number 2^8 is known in computer terms as a byte and consists of 8 bits (256 in binary notation), thus the resolution of a digitised image is described by the number of pixels and the number of brightness levels (0 being black and 256 being white). A typical example being 640 pixels by 480 lines with 8 bits of brightness resolution for each pixel. If one byte were used to store the brightness values for each of the three main colours (red, green and blue) then the image would be a full colour 24 bit image (3 x 1 byte) capable of reproducing approximately 16.7 million colours (i.e. $2^8 \times 2^8 \times 2^8$).

One advantage of digital data is that it has discrete values and therefore can be handled without error or corruption. A digital image is therefore adaptable to manipulation by mathematical and statistical functions which can alter the grey level values and the spatial co-ordinates of the pixels.

2.2. THE SYSTEM

A general purpose processing system should be capable of four major procedures: (1) acquisition, (2) storage, (3) processing and (4) display.

The hardware associated with these procedures may consist of :
1. **Acquisition** - video camera, scanner, frame grabber board (digitiser).
2. **Storage** - Magnetic disks (floppy disks, hard disk drives), Optical disks (CD-ROM etc.), Magneto-optical disks. Temporary storage may be achieved in the host computers memory or on suitable frame grabber boards
3. **Processing** - host computer, combined frame grabber/processor board, dedicated processing board.
4. **Display** - Monitors (monochrome or colour), hard copy devices (laser printers, thermal wax, photography).

2.2.1. Image acquisition

The acquisition of a digital image may be viewed as a two fold process. The first is concerned with converting the required band of electromagnetic radiation (i.e. visible, infra-red or x-ray) into an electrical signal proportional to that of the energy. This is normal achieved using a video camera or scanner. The second converts the electrical signal into a digital format which can be processed by a computer. An analogue to digital converter is usually incorporated into the frame grabber board.

Typical devices used in image acquisition are video cameras and scanners. Video cameras produced today are based on charged-coupled devices (CCD). These are replacing the older Vidicon tube cameras principally due to their ruggedness and low light level operation. These solid-state imaging sensors within a CCD camera consist of an array of discrete "photosites' each one having a voltage output proportional to the intensity of the incident light. The arrays are arranged in a matrix format with the number of elements present determining the output resolution. This can range from 256 x 256 elements to 2048 x 2048 elements for higher resolution work, with the present 'norm' being either 640 x 480 or 768 x 512.

The scanner uses the same principals as the above mentioned camera with the exception that the 'photosites' are arranged in a row and the image is produced by the relative motion of the detector and the object.

The video output of both devices must then be fed into a digitiser which quantizes the signal producing the digital output. The digitiser is normally accommodated on a specifically manufactured frame grabber board which slots into the computer. A basic frame grabber board consists of an input digitising section (an analogue to digital converter - ADC), at least one frame buffer (a memory store for the digitised image) and an output section for converting the digitised signal back into an analogue signal for display on a monitor (a digital to analogue converter - DAC). Invariably the frame grabber board will have additional features such as additional frame stores, control for graphics overlay and in the more expensive boards provision for undertaking a number of the standard image processing routines such as histogram equalisation, convolution and even Fourier Transforms.

2.2.2. Image storage

Digital images require large amounts of storage space. The size of the image file is measured in bytes and is related to the number of pixels in the image. For example an 8-bit digital image measuring 1024 x 1024 pixels requires one mega byte of storage whilst an 8-bit digital image measuring 640 x 480 pixels i.e. one that will fill a standard computer monitor requires approximately 300 Kbytes of storage space.

The storage of images may be viewed at a number of levels - temporary storage, short term storage and long term or archival storage. The type of storage required dictates the most suitable storage medium.

2.2.2.1. Temporary storage

Temporary storage is used whilst acquiring images, for the fast recall of images and during the processing of images. Provision of temporary storage may be through the use of computer memory or through the 'frame-buffers' incorporated into frame grabber processing boards. The latter method has the advantage in that the images may be accessed very quickly whilst freeing the computers memory for processing functions. The images held within these systems will be lost when the computer system is switched off.

2.2.2.2. Short-term storage

Magnetic disks such as the computer hard disk or the floppy disk provides a convenient, permanent means of storing digital images in the short term. The hard disk system provides the most convenient and faster recall time however it should be viewed as primarily storage for programs and often used data as a disk nearing its capacity can effect the computers speed and efficiency. The standard 3½ inch floppy disk has a much slower access time and could not be considered usable during an image processing session. One solution to this problem is to temporarily load the required images from the floppy disk onto the computers hard disk before the processing session and once finished transfer the data back. A further disadvantage of the floppy disk is their restricted capacity. A 1.44Mbyte high density disk will only hold four 768 x 512 images.

2.2.2.3. Long term storage

This is a fast growing area of disk technology with Optical and Magneto-optical devices providing a constantly developing range of high capacity storage. Disks and drives are available which provide read-only facilities, write-once-read-many or erasable-rewriteable media with capacities of 600Mbytes or more. CD-ROM's typify read-only devices, write-once (not rewriteable) types are termed WORM drives whilst Magneto-optical disks (MO) provide erasable-rewriteable storage. It is presumed that they all will have their part to play once the technologies have become more standardised however at present there are the problems of compatibility and continued product support to consider.

2.2.3. Image processing

The nature of a digital image is such that the procedures used in processing the image may be expressed as mathematical algorithms. Thus most image processing functions may be carried out in software by the host computer and only require dedicated hardware - in the form of additional processing boards with an Arithmetic Logic Units (ALU) - when the processing speed is crucial.

Image processing can be very subjective and is characterised by specific solutions. Thus functions which may work well on one image may provide little or no advantage on another. However there is available a range of proven techniques which provide a starting point for the solution of specific problems. It is therefore essential that available image processing software not only provides the standard techniques but also the means of modifying or customising them to suit the needs of the image and the application.

2.2.4. Display

The principal means of display is through a cathode ray tube (CRT) based computer monitor. Most applications employ a dual monitor configuration - one for the display of images and one for the image processing program interface. The display monitor is driven from the output stage of the frame-grabber board which for monochrome boards consists of three DAC one for each of the red, green and blue channels thus permitting false colouring of the output image. The colour CRT monitor utilises three electron beams (red, green and blue) focused, through the holes in a 'shadow-mask', onto a triangle of phosphor dots, one dedicated to each colour. The intensity of each beam is determined by the applied output signal from the three DAC's. The distance between the dots on the 'shadow-mask' is termed the dot pitch - the smaller the pitch the sharper the image. A 0.28mm dot pitch is currently the most common in high resolution monitors.

2.2.5. Hard copy output

At the present time there is no means of producing low cost, high quality hard copy output from a computer based system either in colour or grey scale. The options available for producing what could be termed medium quality output are considered below. It should be noted that this is an area in which the technology is developing fast and subsequently higher quality will become more widely available and affordable.

2.2.5.1. Laser printers

This discussion will be limited to black and white laser printers for although colour laser printers are available their cost at £15,000 to £25,000 is prohibitively high for the majority of users. Laser printers operate by shining a pulsed laser light onto a spinning polyhedral mirror. The deflected beam scans across a rotating negatively charged drum giving the exposed areas a positive charge. Toner is attracted to the charged areas and from there on to the paper by a charged element placed behind the sheet. Finally, heated rollers fuse the toner to the paper.

The resolution of a laser printer is measured in dots per inch (dpi) with the majority of available printers having a resolution of 300dpi. By altering the amount of toner used and the pattern of dots on the page approximately 60 grey scales levels can be produced providing an acceptable reproduction of the image. Higher resolution printers capable of 600dpi are now becoming available (HP LaserJet 4M and Apple Laserwriter Pro 630) at a cost of £2,00 to £2,500, however this does not give an increase in the number of grey scale levels only in the 'fineness' of the image. An exception is the new Apple Laserwriter Pro 630 utilising a feature termed Apple Photograde which claims to increase the number of grey scales although it only operates at a resolution of 300dpi.

Examples of laser printed output are shown in figures 25 and 26. Both images were produced on a HP LaserJet 4M with figure 25 at a resolution of 300dpi and figure 26 at a resolution of 600dpi.

2.2.5.2. Thermal printers

There are two main types of thermal colour printers, thermal wax and thermal dye sublimation. Thermal printers use heat to transfer colourant from a printing ribbon to the specially prepared paper. The thermal wax technique is only capable of transferring a fixed amount of colourant and must use dithering (the representation of different colour values by different patterned dots) to achieve an acceptable range of colours thus giving a rather speckled appearance. Thermal dye sublimation printers can vary the volume of colourant transferred to the paper in 256 steps, along with the intensity. This produces a continuous tone effect similar to photography however the resolution is often lower than 300dpi producing jagged edges to curves etc. Thermal printers are expensive, costing between £2,000 to £11,000 to buy and from 30p to £4 (dye sublimation printers) per sheet to run.

2.2.5.3. Photography

Photography can be an attractive option for capturing the true appearance of the image on the monitor screen and is often the only means of achieving realistically priced colour output. The equipment need consist of no more than a 35mm camera and a tripod. The process must be carried out in a darkened room to avoid any reflections in the monitor screen and care must be taken in positioning the camera to preserve the image aspect. Better results will be obtained from the larger flatter screens as the effects of the screen curvature are not as obvious. The main consideration in recording images photographically is the use of shutters speeds of longer than 1/25 th of a second (i.e. greater than the refresh rate of the screen) in order to compensate for the fact that the shutter is not synchronised with the monitor control signals. Attention should also be paid to the processing of the film in order to obtain the correct colour balance, particularly if using the fast automated services for colour prints.

Except where indicated, all the images used in this report were recorded photographically using Kodak 100 ASA colour print film. The set up consisted of a Nikon FE 35mm SLR camera mounted on a tripod positioned approximately 80cms from a Sony PVM Trinitron display monitor. Through the lens metering was used to determine the exposure of 1/4 sec. f-8 for a blank screen image of 132 grey scale value. This exposure was used for the majority of the images with some modification if the image was particularly dark or bright. A more accurate quantification of global image brightness and hence the required exposure would be to calculate the average pixel value for the whole image.

A further consideration is the effect that the hard copy device will have in the reproduction of the screen (monitor) image. Typically there may be some compression or expansion of the dynamic range or alternatively the device may not be capable of recording the complete spread of intensities. Further processing may therefore be necessary prior to final output.

3. PROCESSING TECHNIQUES

The main types of processing functions may be conveniently divided for discussion into the following groups:-

1. Geometric processing
2. Pixel point processing
3. Pixel group processing
4. Frequency domain processing

Whilst processing functions within one area may appear to duplicate those from another, there are generally differences in terms of actual pixel values which will then dictate the effects of further processing.

The majority of processing functions involve large numbers of computations which even on the fastest PC's can take time. One characteristic of radiographs is the often large expanse of dark background surrounding the object image, elimination of this, and other areas not of interest from the processing computations can effect a more efficient use of the functions. The selection of these areas or regions of interest (ROI) from within the image and the ability to process these independently of the background is an important consideration in applying image processing functions. In addition to the reduction in computation time, the resulting image is often more readable as background noise can produce visual distractions.

3.1. GEOMETRIC TRANSFORMATIONS

Geometrical transformations are concerned with the scaling and rotation of images. The operations are used frequently in computer graphics and are performed by moving pixels from their original co-ordinates to new co-ordinates in the output image. The use and benefits of the techniques are mainly in the pre- and post - processing areas, for example manipulating the images for selection of the area or region of interest and in the presentation of images.

3.1.1. Rotation

The rotation of an image about a central point can be described by the following equation in which θ is the angle of rotation.

$$x' = x * \cos(\theta) + y * \sin(\theta)$$

$$y' = -x * \sin(\theta) + y * \cos(\theta)$$

The rotation functions implemented in the majority of off-the-shelf processing packages are often restricted to 90°, 180° and 270°. This is due to the fact that the sine and cosine of θ would be fractional, thereby necessitating the use of interpolation to place the calculated pixel co-ordinate to the nearest integer pixel position.

3.1.2. Scaling

Image magnification can be achieved by replacing each pixel from the original frame by an array of pixels, the size of which depends upon the choice of scaling factor. The brightness of these 'new' pixels is taken to be equal to that of the original input pixel.

This method of magnification is not entirely satisfactory, as the arrays of new pixels quickly become visible as rectangular blocks of discrete grey levels. This produces a marked chequer-board effect within the resulting image. There are available a number of algorithms based on fractal geometry which can overcome this effect to some extent. However these do not as yet appear to have filtered through to image processing. The relevance of these techniques must also be questioned as they are viewed as adding information to the image which is not contained in the original.

A more satisfactory solution to magnification is the acquisition of the image at the required size prior to processing. Figures 31 to 34 illustrate the above discussion. The original radiographic image, measuring approximately 2.5cm in diameter (figure 31), benefits from the initial magnification during image capture. Two times magnification of this digital image (figure 32) produces a slight blurring of the fine details whilst at a higher magnification the individual pixels become visible (figure 33). Magnification of the radiograph and subsequent image capture through a low powered microscope is shown in figure 34. The improvement over magnification of the digital image is substantial and illustrates well the large amount of detail radiographic film can record.

Reduction in image size can be achieved by a number of methods the most common however is to compute the average intensity value of the array of pixels to be scale down and use this value for the output pixel. Reduction can prove very useful when comparing and presenting the effects of processing techniques, this is illustrated in figure 35 in which each image has been reduced by 25 %.

3.2. PIXEL POINT PROCESSING

A term which refers to the application of the same processing function to the intensity value of each individual pixel within an image without taking into account neighbouring pixels. Pixel point processing may be divided into operations performed on a single image and those performed on multiple images. Multiple image operations will be discussed at the end of this section as they generally entail the combination of images after more specific processing techniques have been employed.

3.2.1. Single image processes

3.2.2. Intensity transformation

An simple and effective method of altering the brightness and contrast of an image is through the use of LUT's (Look Up Tables). The look-up table holds a range of intensity values which can be used to re-map the intensity values of the original image, thus for example changing the original grey-scale which would typically be a linear progression from 0 to 255 in single unit steps to one using a logarithmic or exponential progression from 0 to 255. The means of expressing this re-mapping operation is typically in graphical form with the abscissa being the original intensity value and the ordinate the new intensity value (figure 6). Look-up tables can be used for stretching or compressing the dynamic range of an image, for high-lighting specific grey-levels and for adding false colour to an image (see separate section on false colour). The advantages of using LUT's are ones of speed; once a LUT has been loaded into memory, the image processing is in real-time, and flexibility, any number of LUT's may be stored on disk, called up when needed or modified for a specific image. The main disadvantage is that the effect is global i.e. the same operation is applied to the whole image. The following examples illustrate some of the uses of LUT's.

3.2.2.1. Image inversion

Inversion of the intensity values within an image i.e. 0 becomes 255 and 255 becomes 0, can often reveal details in the lower range of grey scale values. This may be attributed to the non-linearity of the human visual system to image intensity. Figure 7 shows a plot of the inverted LUT with the original image and the resulting processed image in figure 37 - the original image may be seen in figure 36. More detail is revealed of the outer corrosion products and the perimeter of the image.

3.2.2.2. Expansion of the dynamic range

Low contrast images are characterised by a small dynamic intensity range i.e. the majority of the pixels are clustered within a narrow range of grey levels. By expanding the range of intensity values the contrast can be increased producing a more visually informative image. This can be achieved by using a LUT based on a function such as the one illustrated in figure 8. The original range of values e.g. 10 to 70 on the x-axis is expanded to 10 to 150 on the y-axis. The example illustrated in figure 38 shows greater detail within the outer corrosion layers this was achieved by using the function in figure 8. This technique can also be applied to discrete ranges of intensity values in order to increase the contrast and therefore the readability of a particular segment of the image.

3.2.2.3. Compression of the dynamic range

This technique can be considered as the opposite of the previous example i.e. a large intensity range in the original image, the x-axis is compressed into a smaller range on the y-axis. The technique can be of use in the 'real-time' viewing and 'capture' of radiographs which display a large dynamic range.

3.2.2.4. Grey- level emphasis

Sometimes termed grey level slicing the technique may be used to highlight one or a number of discrete grey level ranges. Two approaches may be considered in the implementation of this technique. The first approach highlights the desired grey level whilst removing the background. In effect producing a binary image (figure 18). The second process highlights the desired intensity range whilst preserving the background (figure 19). The technique may be of use in following boundaries of objects or highlighting flaws within a material.

3.2.3. Histogram modification

A concept which is central to the understanding of digital image processing and in particular point processing is that of the histogram of the pixel intensities over the image or area of interest. The histogram of a digitised image can be expressed as a function showing for each intensity (grey) level, the number of pixels with that specific value. The shape of the histogram provides information on the brightness and the contrast of the image that help determine which particular type of grey scale transformation may be appropriate. In some respects it may be viewed as providing a measure of the 'quality and content' of the image. By altering the position and the shape of the histogram it is therefore possible to alter the intensity characteristics of an image. This may be achieved by a number of methods.

3.2.3.1. Histogram slide

The original histogram is studied and an appropriate constant can be added or subtracted from each pixel intensity thereby 'sliding' the histogram into the desired region of the grey scale. The operation does not alter the relationship between individual pixels i.e. the resolution of the image, it provides a means of uniformly darkening (subtraction of a constant) or brightening (addition of a constant) the image. Typical values for the constant would be in the range of 25 to 80. Figure 22 shows the histogram of the original grey levels of the arrow in figure 39, whilst figure 23 shows the histogram after subtraction of 80.

3.2.3.2. Histogram stretching

This technique stretches or shrinks the original histogram of an image. This is achieved by either the multiplication (stretching) or division (shrinking) of each pixel intensity by a constant. Stretching the histogram has the effect of uniformly increasing the contrast of an image in addition to brightening a dark image whilst shrinking the histogram may be used to reduce or compress the contrast and darken a bright or washed-out image. A typical value for the constant would be in the range of 2 to 4. Figure 24 shows the effect of stretching the histogram seen in figure 23 by a factor of 2.

3.2.3.3. Histogram slide and stretch

As the typical grey scale is confined to only 256 values the use of the previous two techniques can easily cause 'clipping' of pixel values at the extreme ends of the scale resulting in the possible loss of some image detail i.e. those intensity values which after processing would be greater than 255 or less than 0 will be 'clipped' to 255 or 0, thereby producing areas of constant white or black. This may sometimes be avoided by using combination of the two techniques after a study of the original histogram values. For example a sliding factor may be needed to subtract a small part of the brightness before stretching the histogram or vice-versa. The effect of this technique is illustrated in figure 40 in which the central area of the arrow shown in figure 39 has been processed. Compare the histograms in figures 22 and 24.

3.2.3.4. Histogram equalisation

The aim of histogram equalisation is to re-assign the intensity values within an image in order to produce a flat histogram stretching from the grey level values of 0 to 255. This is not fully achieved since in practice equalisation is attained by stretching the peaks and shrinking the troughs of the original histogram thus the different intensities are spread more evenly throughout the image. Figure 20 shows the histogram of the coin in figure 26 whilst figure 27 shows the coin after equalisation and the new histogram can be seen in figure 21. The technique is widely used and has the effect of increasing the dynamic range of the image thereby presenting the data in a better form. Graininess may result from images which originally had a narrow range of intensity values and areas of poor digitisation may become more apparent.

3.2.3.5. Local histogram equalisation

This function could be viewed as a pixel point or a pixel group operation depending upon whether a single pixel or a group of pixels receive the resulting value. It is included here for reasons of association with the histogram equalisation function.

Although histogram equalisation as described above can be very effective it is a global technique and therefore does not consider the characteristics of discrete regions. Local histogram equalisation is designed to overcome this problem by calculating a transformation function for each discrete area. A square neighbourhood is defined within the image (e.g. 15 x 15 pixels). The histogram of this area is computed and a histogram equalisation function is obtained. The function is used to map either the central pixel or, alternatively all the pixels within that area. The neighbourhood is then moved to an adjacent pixel and the process repeated. Altering the size of the neighbourhood produces variations in the visual result. An example of this technique is shown in figure 43 with the original image in figure 42. Local enhancement may also be based on other properties exhibited by the image, for example the mean of the grey levels (the brightness) and the variance (a measure of the contrast). Local enhancement by histogram equalisation can be a time consuming process (10 to 15 minutes) and is therefore not normally included as a feature of 'off the shelf' image processing packages.

3.2.4. Pseudo colour techniques

The use of colour in image processing is motivated by the fact that the human visual system is able to discern a vast range of colour shades and intensities compared to approximately 60 to 90 JND (just noticeable differences) in a grey scale (JND is a measure of perceived changes). Pseudo colour processing assigns a shade of colour to a grey scale level or range of intensities, the aim being to facilitate the identification and extraction of features from the image.

The process is carried out using three LUT's, one for each colour component i.e. red, green and blue, to transform the original grey levels into colour variations. The range of colours depends upon the selected transformation functions for each of the three colour components. As is described in section ?? these functions can be represented graphically however in this case the original grey level is plotted against the new colour component value. This requires three plots one each for the red, green and blue elements.

The requirements of an effective colour scale are that it should preserve the perceived order of values within the image, it should convey the distance between the values the colours represent and that they should not create perceived boundaries that do not exist within the original image.

A typical example of a pseudo colour transformation can be seen in figure 16 with the resulting image in figure 45. This colour scale uses a similar convention to colour contouring of maps which assigns green to low values (low ground), through yellow and brown to mauve or white for high values (high ground). The order of values within the image is preserved whilst the changes in the grey scale 'landscape' are more easily recognised. Other commonly used colour scales are the heated object scale and the magenta scale. The heated object scale (shown in figure 10 with an example in figure 48) runs from black to red (lowest grey level values), through orange and yellow to white (highest grey level values) and is based on the fact that the human visual system has maximum sensitivity to luminance changes for the orange-yellow hue (compare the vast range of these shades in the Munsell colour system). The magenta scale (shown in figure

14 with an example in figure 49) is based on the fact that the human visual system can be most sensitive to hue changes within the magenta range.

All the intensity transformation functions discussed in section 3.2.2 are applicable to pseudo colouring however there is the potential for a greater sensitivity to grey level change due to the possibility of being able to combine up to three different functions (one each for the red, green and blue components) within a single transformation.

Although a simple technique to implement the addition of colour is a very powerful tool with which to highlight and enhance images for human perception.

3.3. PIXEL GROUP PROCESSING

This area of image enhancement is concerned with the modification of the intensity of a single pixel whilst taking into account the intensity values of the neighbouring pixels. The two main types of operation are firstly the convolution of an area of the image (normally 3 x 3 pixels) with a mask or kernel of coefficients, and secondly, mathematical operations performed on the ranked intensity values within an area of the image (typically a 3 x 3 pixel square or a vertical cross of 5 pixels).

3.3.1. Convolution

The term convolution as used in image processing describes the series of multiplication's and additions between an area within the image (3 x 3 pixels) and a similar sized kernel of values. The resulting value then replaces that of the original central pixel. The operation multiplies each pixel within the defined area with the corresponding value within the convolution kernel and sums the products.

The values within the convolution kernel dictate the type of processes performed on the image. These may include high and low pass filtering, edge detection or high-lighting and texture representation.

3.3.2. Spatial filtering

The spatial frequency within an image can be regarded as the rate of change of intensity expressed as a function of distance. Hence large changes in intensity occurring in close proximity to one another represent high spatial frequencies (e.g. edges, fine details and noise) whereas low spatial frequencies are characterised by regions with small changes in intensity. In filtering out the high or the low frequencies it is possible to remove noise or small details from an image, to sharpen edges or highlight fine detail.

3.3.2.1. High-pass filter

A high-pass filter is designed to remove or reduce the low frequency components of an image thus sharpening details or enhancing areas which are blurred. It has the disadvantage of enhancing any noise within the image which may make the end result unacceptable. A typical high-pass filter is shown in fig 1 note the positive coefficient at the centre, negative coefficients on the outside and that the sum of the coefficients is 0. When working in an area with small intensity changes the output of the mask will be very low thereby reducing the low frequency components. Figure 47 shows the effect of using the following filter on the radiograph of the padlock in figure 46.

-1	-1	-1
-1	9	-1
-1	-1	-1

Figure 1 *Convolution mask for a high-pass spatial filter.*

3.3.2.2. Low-pass filter

Low-pass filtering has the effect of blurring the image thereby removing noise and fine detail resulting in an emphasis of the main elements within an image. A low-pass filter is designed to remove the high frequency components from an image. This may be easily accomplished by the use of a mask which will compute the average of all the pixels within the pre-defined neighbourhood. A typical mask is shown in figure 2 below.

1	1	1
1	1	1
1	1	1

Figure 2 *Convolution mask for a low-pass spatial filter.*

3.3.3. Edge detection/enhancement techniques

An edge within an image is considered to be the boundary between regions of distinct intensity characteristics such that the transition can be determined by consideration of the grey levels alone. If the pixel intensity values are viewed on a vertical scale the difference between the values of neighbouring pixels will produce a constantly changing slope or gradient across the image. The greater the difference in neighbouring values then the greater the slope or gradient. Edge enhancement techniques are designed to react to the slope of the pixel intensity occurring within a pixel group thereby accentuating edge details within an image.

3.3.3.1. Gradient operators

The gradient filters operate by accentuating the high frequency components of the image (edges etc.) while attenuating the low frequency components. With these type of gradient operators the enhancement is directional and may be selected from one of eight directions depending on the positioning of the coefficients within the convolution mask. Figure 50 shows the typical effect of applying a gradient filter - this example works in a easterly direction.

1	1	1		1	1	1		-1	1	1		-1	-1	1
1	-2	1		-1	-2	1		-1	-2	1		-1	-2	1
-1	-1	-1		-1	-1	1		-1	1	1		1	1	1

Figure 3 *Convolution masks for directional gradient filters. North, north-east, east and south-east.*

3.3.3.2. Laplacian

The Laplacian edge enhancement operation possesses an omni-directional quality accentuating edge details within an image regardless of orientation. The basic requirement of such a filter is that the coefficient at the centre of the mask be positive whilst the outer coefficients be negative. If the sum of the coefficients is zero then image areas with similar intensity values will produce a zero or low output value. The opposite is true if the input pixel has a very different value to its neighbours and hence the Laplacian filter can be very sensitive to noise. Figure 51 shows the effect of applying a Laplacian filter - in addition to being sensitive to noise, areas of poor digitisation are emphasised.

-1	-1	-1		0	-1	0
-1	8	-1		-1	4	-1
-1	-1	-1		0	-1	0

Figure 4 *Two examples of Laplacian operators.*

3.3.3.3. Sobel operator

The Sobel operator uses two convolution masks to determine the existence of edges in both the X and Y directions. In summing the output of the two masks the resulting edge detection is omni-directional. In addition the process produces a smoothing of the image thus reducing the effects of random noise and is therefore a particularly attractive feature of the Sobel operators. An example of the effect of the Sobel filter is illustrated in figure 52.

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

Figure 5 *The two convolution masks making up the Sobel operators. One detects edges in a horizontal directional, the other in a vertical direction. By combining the results a omni directional mask is effected.*

3.3.4. Non-Linear Operators

As with the convolution process described above, non-linear operators act on neighbourhoods. The operations are however based directly on the values of the pixels under consideration rather than on the use of coefficients. Such operations may compute the median value within a neighbourhood or react to the maximum or minimum values.

3.3.4.1. Median

A median filter is an alternative to a low-pass filter for the reduction of noise without significantly blurring the image. As described above it is a non-linear filter in which the pixel values within the pre-defined neighbourhood are ranked, the median value determined (rather than the average) and this value is assigned to the central pixel. This has the effect of eliminating any isolated intensity points from within the area.

3.3.4.2. Edge enhancement

Edge enhancement using non-linear operators may be accomplished in a number of different ways. An operation which highlights edges whilst negating the background computes the difference between the maximum and the minimum intensity values within the neighbourhood and places this value at the central pixel. Areas of an image which produce a great range of values are therefore emphasised whilst those with little or no variation are dark.

A more subtle type of edge enhancement may be achieved by comparing the difference between the maximum and minimum with a threshold value. If the value exceeds the threshold then the pixel to be processed is set to either the maximum or the minimum value which ever is the nearest to the original, otherwise the median value is selected.

3.4. FREQUENCY DOMAIN PROCESSING

The previous discussion has concentrated on processing images within the spatial domain i.e. the two dimensional position of each pixel and its grey level have been accessed directly. An alternative means of examining an image is in terms of the spatial frequency of its components. Mathematical analysis can show that any wave shape can be approximated by a sum of periodic functions (i.e. sine and cosine functions). As the number of functions increases so the resulting wave form more closely represents the original spectrum. The spatial frequency was defined earlier in section 3.3.2 as the rate of change of intensity as a function of distance. A graphical representation of this data can therefore be approximated by a series of sines and cosines of various frequencies. Each individual sine or cosine wave of a specific frequency is therefore a component part of the image. The most common mathematical function used to achieve this conversion is termed the Fast Fourier Transform (FFT). Precise manipulation of the image is therefore possible by applying functions which will act upon the frequency components. The most common application is in removing specific frequency bands (normally high or low) representing periodic features within the image e.g. interference of specific frequency. Plotting the square of the amplitude of the frequencies gives an energy spectrum thereby providing a graphical representation of the frequency power and distribution (figure 58). This can be used for selecting the desired frequencies for manipulation. After manipulation of the frequency components an inverse Fourier Transform is applied to convert the data back into the spatial domain.

3.4.0.1. High-pass filter

The frequency data is multiplied by a weighting factor to reduce the low frequency components of the image thereby sharpening edges and features. Examples are shown in figures 44 and 54.

3.4.0.2. Low-pass filter

A low pass filter reduces the high frequency components and can be useful for controlling periodic noise within the image.

Frequency domain processing is more applicable to problems with a global or periodic structure where as spatial domain processing favours problems with local structure i.e. non repeating features.

3.5. MULTIPLE IMAGE OPERATIONS

Multi-image operations are used to combine two or more images. The pixel values of one image are either added or subtracted from the corresponding values of a second or subsequent image. The process is often used for the combination of a processed image with the original image to aid interpretation of enhanced features or to provide a more 'natural' looking result.

3.5.0.1. Image subtraction

Image subtraction can be used to show the difference between frames and for the removal of background shading etc. An example of this technique which may remove the typical 'fogginess' from an archaeological radiograph is the subtraction, from the original, of an image processed with a low-pass filter (see section 3.3.2.2 for details). The result highlights the high frequency components within the image i.e. those associated with edges and detail.

3.5.0.2. Image addition

The addition operation has the effect of superimposing one image on another. An example would be the addition of an original image with one in which feature edges had been detected using the Laplacian filter (section 3.3.3.2). The resulting image can have a naturally sharpened appearance whilst retaining all the details and subtle variations of the original.

3.5.0.3. Frame averaging

This technique is normally used during the initial capturing of digital images. Noise is inherent in most electronic instrumentation. The lower the original signal level and the greater the amplification then the greater the level of noise. When capturing an image for digitisation the lower the light level the greater the noise produced from camera. As this is generally a random phenomena capturing a series of images of the same scene and computing the average of the corresponding pixels will produce a significantly enhanced image. The most common type of averaging technique is termed recursive filtering. It uses a weighted averaging technique to produce the final image from a sequence of images of the same scene. The advantage of this technique is that it does not require the whole sequence of images to be stored as do similar techniques such as modal filtering and median filtering.

3.6. DISCUSSION

This section has outlined the underlying concepts and the basic procedures currently in use in digital image enhancement. As mentioned previously, image enhancement is characterised by specific solutions to problems,the results often being very subjective in nature. Whilst some of the standard techniques can work well over a range of images (histogram equalisation is a good example),it is more often the case that the best solutions are offered by a combination or customisation of techniques. Some of the more robust combined operations have been mentioned and a description of their intended objectives has been offered. In addition it should be noted that some of the more advanced techniques described above may not, in isolation,produce visually informative or satisfying images. Frequently there is a need to alter the brightness and contrast of the image using,for example,the histogram slide or stretch techniques in order to better evaluate the resulting image. Customisation of techniques often takes the form of modifications to the coefficients in a convolution kernel, re-scaling the results of a convolution or, in the case of a Fourier Transform for example, which particular frequency bands need to be accentuated or attenuated.

For every application however there could be a number of possible solutions of varying degrees of complexity,with the only judge of how well an particular method works being the viewer. It is therefore necessary that the control of the image enhancement process be, in this particular instance,the responsibility of the conservator.

4. AN APPROACH TO THE IMAGE PROCESSING OF RADIOGRAPHS

The techniques discussed in the previous section provide the conservator with the means of manipulating images in order to visually improve the image, extract specific information and facilitate the interpretation of the image details. As stated previously image processing is characterised by specific solutions to specific problems. This applies particularly to archaeological materials with the non standard nature of the artefacts and the encountered variations in physical and chemical conditions. It is therefore not possible to recommend techniques or processes which will solve all imaging problems within conservation radiography however it has become clear that the adoption of a structured approach towards image processing can greatly improve the quality and the quantity of information gathered from radiographs. The approach which has been established by this study may be divided into the following three sections (a) the identification of a series of techniques which will generally improve the quality and 'readability' of the radiograph and therefore make subsequent processing more effective; (b) the identification of techniques which will solve specific, frequently encountered problems associated with conservation radiographs; and (c) to highlight those processes which may afford some measure of success in enhancing features often when applied in combination with other image processing techniques.

4.1. STAGE ONE - PRODUCTION OF A HIGH QUALITY DIGITAL IMAGE

This area is concerned with the initial digitisation and processing of the radiograph in order to achieve the highest quality image. The areas considered are those of image capture - including lighting conditions and frame averaging, image scaling, treatment of the background areas and contrast adjustment via LUT's and histogram stretch and slide. The outcome of this first processing stage is the production of an image with a broad, evenly distributed contrast range, a suitable brightness level for viewing, minimum interference from noise and a constant grey scale valued background.

4.1.1. Image Capture

Lighting conditions and magnification should be optimised to provide the best possible representation of the image for digitisation. The recommended conditions for viewing x-radiographs apply for image capture. A quality X-ray viewer will provide even background light levels and screens should be used to mask any surrounding light. If the image or the section of interest is very dark then an appropriate high intensity viewer should be used. The camera should be positioned perpendicular to the plane of the radiograph and at an appropriate distance to provide as large as possible a representation on screen of the desired image. The light level should be adjusted either from the X-ray viewer or from the camera/lens to avoid any under or over exposed regions within the area of interest. This is important in avoiding loss of grey level detail by digitising large areas to either 0 (black) or 255 (white). If this is unavoidable as with an image of very high contrast then the image may have to be split into sections and reassembled after digitisation.

Electronic noise inherent within the image capturing system gives rise to image degradation and spurious features within the image. Processing can often enhance this noise and it is therefore desirable to reduce it to as low a level as possible. The use of correct lighting conditions and exposure play an important part in the reduction of noise however further reduction can be achieved by using the frame averaging technique described in detail in section 3.5.0.3. In general to produce a digitised image with as low a noise level as possible an average of corresponding pixels from at least 10 - 15 images should be computed.

4.1.2. Image scaling

The use of suitable magnification at this stage is also important in recording all the required detail. The magnification of a digital image is very easily accomplished however with an image resolution of 768 x 512 pixels it does not produce satisfactory results as the individual pixels become visible. A fine grain x-ray film however is typically composed of 0.25 micron grains and can be magnified up to 50 times before the individual grains become visible. Magnification before digitisation is therefore preferable unless operating at a very high resolution (see example in section 3.1.2 and refer to figures 31 to 34

4.1.3. Treatment of the image background

Advantages may be gained from processing discrete regions of interest rather than the whole image. Due to the nature of radiographs of archaeological materials, there is generally a large expanse of black surrounding the artefact. After digitisation these areas are normally within the grey scale range of 0 to 12. Including these areas in any histogram computation and/or transformation process will provide neither accurate information on the grey scale characteristics of the object image nor an appropriate transformation function. The resulting histogram would be weighted towards the zero or dark end of the scale. The ability to define a region of interest (of any shape) on which to operate is therefore of great importance. The term 'any shape' is stressed here as regions of interest which can only be rectangular or circular in shape are not always appropriate for archaeological materials. When using the more complex processing functions operating over a region of interest can radically reduce the processing time. The processing of a background which is not of an absolutely constant grey scale value can produce very unsatisfactory results negating any improvement within the object image.

An alternative approach of tackling the dark background of a radiograph is to identify the grey scale range present within that area (typically 0 to 12) and use an intensity transformation function to threshold this range to a constant value. The region of interest approach would still be necessary for histogram computations however there are advantages when using convolution procedures as the majority will return a result of zero from areas of constant intensity values. In addition, although zero may typically be chosen for the background value the choice of another grey level - a mid-grey (127) for instance - allows the image to be viewed in a slightly different context. Figure 9 shows a typical LUT for applying an example of this type of background mask with the resulting image in figure 39. Whether this has any advantages is entirely subjective however as with viewing radiographs in negative form it is possible that it is easier for the human visual system is able to assimilate the information available.

4.1.4. Contrast adjustment

Computing the histogram of the image or the area of interest can display some important characteristics that aid in the determination of which particular grey scale transformations may be appropriate. The histogram also offers a measure of the 'quality' of the radiograph and therefore whether the original can be improved upon. Guided by the histogram, appropriate LUT's (section 3.2.2) can be used to re-map the original grey scale values, compensating for the deficiencies in the original. Alternatively the histogram stretch and slide technique (section 3.2.3.3) may be used to adjust the grey scale range and distribution within the image.

Histogram equalisation could be considered in this section as being a technique which improves the overall quality of an image, however, it can produce extremes of contrast, resulting in a harsh unrealistic effect which often does not lend itself to further processing.

4.2. STAGE TWO -

The second stage in this approach to image processing follows on directly from the initial production of the quality image and tackles the more general problems which restrict the 'readability' of the radiograph. One characteristic which is often apparent and can restrict the interpretation and obscure detail within a radiograph is the 'fogginess' of the image. This can be attributed to the small difference in x-ray density between the object and the often large volume of surrounding corrosion and the unsharpness caused by the depth of the object. Radiographic film is also capable of recording a broader range of image densities than is discernible by the human visual system at any particular instant. Image details may therefore be recorded which are not immediately apparent to the observer.

The aim of this section is to reduce the 'fogginess' of the radiograph thus presenting a sharper, 'cleaner' image for interpretation and to scan the image for hidden details which may not be immediately detectable by the eye. In addition there is often a need to specifically highlight certain features in order to draw attention to them and facilitate their interpretation.

4.2.1. Intensity transformation

The application of an intensity transformation via the use of LUT's will provide 'real time' processing of the image to either expand the lower grey scale values or produce a negative image. This will make visible any detail within the darker areas of the image (typically within the range 10 to 50 grey scale value) providing information on outer corrosion products and other characteristic areas of low X-ray absorbency. This can aid in revealing the 'true' edges of objects within corrosion layers (refer to section 3.2.2 and figures 8 and 37).

4.2.2. Histogram equalisation

Histogram equalisation is a very useful and widely applicable technique. Vague image characteristics may become very apparent after application of the equalisation process to the whole or a section of the object image. The technique is particularly useful for artefacts which exhibit a wide range of thickness and density, the extremes of the image are equalised giving an equal emphasis to all areas within the image. Fine details can become more apparent and the tracing of features through areas of differing densities is facilitated. In many instances histogram equalisation is a useful technique for indicating whether any hidden detail and information exists within the radiograph and as such very little difference may be seen however when processing the whole object image of a 'good quality' radiograph (refer to section 3.2.3.4 and figures 28 and 29).

4.2.3. Local histogram equalisation

This technique can reveal all available detail within an image, however, as the process applies an equal weighting to all features, the initial result is often very confusing and unreadable. There is generally a need for further processing to improve the 'readability' of the image and as such this will be dealt with in more detail in the following section.

4.2.4. Pseudo colour LUT's

Pseudo colour LUT's may be used in the context of this section for two purposes. The first is the recovery of detail not immediately apparent to the observer and the second aims to highlight visible detail to facilitate interpretation and for the presentation of data. The principles behind the use of pseudo colour are dealt with in detail in section 3.2.4 and are based on the ability of the human visual system to recognise a greater range of colour shades and intensities than those within a grey scale.

Revealing hidden details through the use of colour involves experimentation in scanning the image with different LUT's covering the complete grey scale range from 0 to 255. Those LUT's based on the high sensitivity of the human visual system to luminance (orange/yellow) and hue (magenta) are those which theoretically should be the most successful and therefore provide the suggested starting point for further experimentation (see figures 10 and 14).

In order to highlight detail which is already visible the grey scale range of the detail is identified and a colour LUT created which will assign the required colour to that range. Figure 11 shows an example of a specifically created LUT with the resulting image in figure 55 - note that the colours used to highlight the area run from dark red to red, this appears to be more effective than choosing a single colour for the whole range. This technique can be very effective in following features such as plating or inlay across an image. Once the most suitable LUT has been applied to the image editing may be required to remove areas with the same grey scale range to further clarify the required details. With certain software, such as that produced by Foster Findlay Associates, it is then possible to extract the highlighted area and save it as a separate image. This feature can be of use when comparing details from a number of images.

4.2.5. Image sharpening

Sharpening of the image can be accomplished using high frequency filters such as described in section 3.3.2.1. These will emphasise any sudden changes in pixel intensity across the image resulting in sharper edges to features. The technique appears to work best as a final processing function rather than as one for preparing the image for further treatment. The quality of the resulting image is dependant upon the typeof material from which the radiograph was taken. Copper alloys, for example produce very grainy images presumably emphasising the variations within the internal structure of the alloy (figure 56).

Care must be taken when sharpening an image as any noise or random high frequency areas will be emphasised producing unsatisfactory results. The application of a smoothing function such as the median filter before sharpening the image may reduce the effects of noise within the image.

An alternative approach to image sharpening is through the use of FFT's, converting the image into the frequency domain, applying a high pass filter and subsequent conversion back to the spatial domain. This is a slightly lengthier process than the previous and requires some experimentation in order to select the correct frequency cut off point for the particular subject. The image produced appears clear of much of the 'fogginess' of the original and whilst the edges of features are not as sharp and well defined as the previous treatment there is less of a processed look about the image (figures 44 and 54 illustrate these points).

4.3. STAGE THREE -

Experience has shown that in general the simpler the processing function used the more widespread and successful is its application. The more advanced techniques tend to require care and precision in their application and appear to be more useful for 'fine tuning' features within an image or for solving very specific problems. The previous section dealt with the more widely applicable processing functions enhancing the image to reveal and/or clarify details within the radiograph. The third and final stage deals with functions which generally need to be used in combination with others to produce the required results.

Radiographic images of archaeological artefacts are often very complex being formed by the superimposition of all surfaces and internal features (at all depths) of a three dimensional object into two dimensional space. In the previous section the image was processed in order to reveal all the recorded detail however this can often result in the production of too much information further compounding the problems of interpretation or presentation. In these cases a solution would be the selective enhancement of relevant detail and/or the removal of unnecessary features thereby improving the 'readability' of the radiograph and effecting a more positive interpretation.

Due to the nature of the radiographic image, the individual requirements of the observer and the nature of the processing functions it is not possible to suggest specific function combinations with which to accomplish specific tasks. It is however possible to record some observations on the use of various function combinations. This stage of processing is, by necessity, one of experimentation and as such relies on the knowledge and experience of the conservator in radiographic interpretation and also in having an understanding of the mathematical concepts behind the image processing functions.

4.3.1. Edge detection and enhancement

In theory, edge detection or enhancement appears to be very useful procedure. In practice the great amount of detail held within a radiograph produces an output which is often more difficult to interpret than the original. Ideally some form of selective edge enhancement or pre-processing of the image is necessary in order to obtain an effective result. One method which has been suggested earlier (section 3.5.0.2) is that of combining the original image with the edge enhanced version. This can provide some measure of success in addition to sharpening the details within the original image.

Some of the functions used for edge detection, for example the Sobel filter, operate more successfully on larger features exhibiting steep grey scale gradients. The fine detail and subtle variations in archaeological radiographs do not appear to respond well to these filters (often the image goes very dark), however, by applying a suitable histogram stretching function, the fine edge details can be very successfully enhanced revealing a great deal of information. The resulting image can then be combined with the original or an equalised image for a more selective enhancement effect. An example of this process is illustrated in figure 52.

4.3.2. Local histogram equalisation

The theory behind local histogram equalisation was discussed in detail in section ??. Generally 'off the shelf' software does not include this type of function as the process can be time consuming and some of the parameters can be image specific. The technique however can produce some very good results as shown in the work of Watkinson and Lange (1992) and in figure 43 in this study. Accompanying this report therefore is the program ARCH_X which includes a number of variations of the local histogram equalisation technique. The functions are used to their best advantage on images which have undergone the stage one processes. It is important that the background is a constant grey scale value, as, if not, the effect seriously detracts from any information revealed within the image. As discussed in section 4.2.3, the equalised image can be very confusing with equal emphasis being given to all details. Further processing is necessary and this may take the form of successively combining the equalised image with the original to negate some of the detail or adding or subtracting an edge detected image in order to highlight features. In addition the effect may be softened by the use of smoothing filters and/or adjustment of the contrast range. An intriguing possibility is confining the application of the technique to a specific grey scale range, thereby preserving sections of the original image whilst equalising the remainder.

4.3.3. Texture

Gradient filters are constructed in order to accentuate the high frequency components within an image i.e. the large changes in grey scale values as typified by 'edge' features. The filters are also directional and the combination of these properties together with additional contrast adjustments may be used to give a textured surface appearance based upon the variations in grey scale values across the image. An effect similar to that of illumination by low raking light may be achieved by application of the correct convolution mask (this will also define the direction of the effect) and subsequent adjustment of the contrast and/or brightness. The choice of convolution mask requires some experimentation to suit the particular image as does the type of additional processing. An example of this technique is shown in figure 57 - the filter used works in a south easterly direction. The combination of textured and original image may have certain advantages in varying the strength of the effect.

As discussed previously, this section is very much one of experimentation being image specific and driven by the demands of the conservator. The above techniques are cited as examples of the use of multiple function techniques in order to provide a base from which further experimentation can continue.

This section has formulated a structured approach to the problems of image processing which provides not only a series of guide lines for establishing a 'good practice' in the treatment of radiographic images but also a suitable framework from which further work can be developed.

5. CONCLUSION

As stated previously radiographs are used extensively throughout conservation generally for providing a relatively rapid means of recording the form and condition of artefacts. The high quality and resolution of radiographic film provides a medium capable of recording a large amount of detailed information. However the complex nature of archaeological materials i.e. the extent of degradation and the variation in composition, and the often non ideal exposure conditions can hinder interpretation of this information. A means of improving the quality of radiographs and visually enhancing features within the image would provide more detailed information and aid in their interpretation. This study has shown that digital image processing techniques have much to offer in this respect providing a method for the rapid processing and enhancement of images directly under the control of the conservator.

In order to achieve the best possible results from the manipulation of digital images, it is essential not only to have an understanding of the nature and the quality of the digital image (e.g. the resolution, the grey scale range etc.) but also to have some knowledge of how the different processing functions operate in mathematical as well as visual terms. Considering the subjective results of image enhancement and the vague definition of what constitutes a 'good' image, it is essential therefore that the image processing procedure is under the control of the conservator and that it remains an interactive process. The conservator must be able to define the features that require emphasis and therefore have a good understanding of the techniques of radiography and the interpretation of radiographic images. This will provide the conservator with a sound base from which to make reasoned judgements as to the relevance of the different processing algorithms and hence aid in the interpretation of the results of image enhancement of radiographs.

The digital image and the component parts of the imaging system were considered in section two. From this was established what was considered to be the most appropriate working resolution for this study given the degree of detail required and the cost. Throughout the study the effect of the working resolution was assessed and found, with the exception of image scaling, to be acceptable for all purposes. Whilst high resolution is desirable in many respects e.g. for image archiving etc. and will soon be more widely available, much can be successfully achieved with a resolution 768 by 512 pixels. Similarly, full colour processing systems will soon be available at a more reasonable cost, however, there is a question as to whether these would be appropriate for radiography which is essentially a grey scale image. It is considered that of the system components that are available at present the hard copy device is the least satisfactory. Whilst reasonably high quality printout is available at a price they do not reflect the resolution and quality available from the monitor.

In section three the standard processing techniques most widely available on commercial software was explored and the mathematical algorithms upon which they are based were considered. In addition the visual effect of these processing functions upon radiographic images was assessed. A number of processing techniques and functions not commonly available were identified as being desirable for the specific enhancement of archaeological radiographs. A range of these functions were tested and a programme module was written incorporating those which were considered to be the most widely applicable - the programme 'ARCH_X' accompanies this report. Further refinement of these techniques continues.

The work in section three identified of the more commonly occurring radiographic imaging problems and further experimentation with multiple function techniques led to the formulation of a three stage approach to the image processing of radiographs. The approach is detailed in section four and is viewed as providing guide lines or protocol for establishing a 'good practice' in the digital image processing of conservation radiographs. This is deemed necessary if further progress is to be made in the research and development of radiographic images. Stage one outlines the conditions and techniques necessary for producing a high quality digital image. This is necessary if further processing techniques are to be successful. In stage two the image is scanned for 'hidden' details and where necessary the grey levels are balanced or equalised for emphasis of the appropriate features. The requirements for improving archaeological radiographs are introduced and techniques for reducing the 'fogginess' are suggested. Stage three takes a more experimental stance and the specific requirements of an image are considered. The section discusses various multi-function processes which have been found to provide a good starting point for the further investigation of the image..

This study has concentrated on aspects of image processing of conservation radiographs. However during the course of the work other image forms frequently encountered in conservation were viewed.

Figure 59 shows a typical infra-red photograph of a wooden ink writing tablet. After passing the image through stage one process in order to clean-up the background and adjust the contrast, a sharpening filter was applied resulting in the image in figure 60. Images such as writing tablets which contain discrete blocks of information (i.e. words) are very conducive to image processing techniques as the blocks can be treated on an individual basis with the most appropriate processing function.

Application of a simple sharpening filter has also proved useful for the enhancement of radiographs displaying iron replaced organic features. The example in figures 61 and 62 shows an iron coffin fitting with iron replaced wood before and after processing.

As mentioned in the previous sections histogram equalisation is a widely applicable function. Figure 28 shows a radiograph of a panel painting the figure being visible superimposed upon the wood. In figure 29 the image has been processed using global histogram equalisation, thus balancing or equalising the grey levels across the image, thereby introducing more detail.

This brief view of some alternative image forms has shown them to respond well to standard image processing techniques. This may be a function of the type of image or alternatively it is due to the fact that the desired result of enhancement is more clearly definable. This reinforces the point that there is still scope for further research and development in the enhancement of radiographs.

The nature of the digital image, i.e. composed of an array of discrete pixels, facilitates the measurement of features within an image. Calculating the number of pixels across a feature or within a feature is an easy task for the computer. With the addition of mathematical formula to the programme rapid calculation of a great number of measurements is possible. In addition the inclusion of a scale digitised at the same magnification as the subject allows the calibration of pixel units to 'real world' units. Some experimentation was undertaken utilising measurement functions on images of textile, fibres and polished sections and the ease with which detailed accurate information can be recovered is impressive. Through the experimental work, a great many uses for this feature became apparent, including the accurate recording of the shape and form of artefacts, the counting of fibres and measurement of degraded textiles and the measurement of inclusions or phases within sections of materials.

Although this study used software supplied by Foster Findlay Associates other products were considered, in particularly those produced by Synoptics and Data Translation. An initial comparison of X-ray images enhanced by these systems revealed that there was negligible difference in the quality of the final image produced and that the algorithm used rather than the software manufacturer was the determining factor in the enhancement of the image. Therefore a detailed comparison between the different systems was not considered relevant. The choice of software for this study was made on the basis of certain operational characteristics. At the time of writing, the main advantages found with the Foster Findlay software was the ability to define regions of interest of any shape (rather than being limited to rectangular areas etc.) and the interactive graphical interface for the creation of LUT's. None of the products reviewed incorporated as standard the more processor-hungry features such as local histogram equalisation. However some form of processing function library was generally available from which, with a little knowledge of computer programming, they could be generated. As the more advanced functions tend to be very application specific, it is clearly more appropriate to develop these 'in house'.

This study is considered purely as an initial assessment of digital image processing as applied to radiographs. It has shown that a wide range of techniques are of great benefit in enhancing radiographs, thereby increasing the recovery of additional information, leading to a more positive interpretation of the form and function of the artefact.

Digital image processing therefore offers great potential within conservation, not only for the enhancement of radiographs, thereby realising further the potential of radiography itself, but also in the rarely considered area of measurement and quantification of radiographic images. In addition it has been shown that its application is not limited to the study of radiographs, thus there is a broad subject area within conservation warranting further research.

BIBLIOGRAPHY

Biek L. 1969 "Artifacts" in **Brothwell D. & Higgs E. (eds)** Science in Archaeology 567-570 Thames & Hudson.

Corfield M. 1982 "Radiography of Archaeological Ironwork" in **Clarke R.W. & Blackshaw S.M. (eds)** Conservation of Iron. National Maritime Museum Monographs and Reports N0. 53

Eastman Kodak 1969 Radiography in Modern Industry. Kodak. New York

Johnson H. 1992 The evaluation of computers in enhancing and analysing X-ray images of archaeological objects. Unpublished BA dissertation. Durham University

Lange E. & Watkinson D. 1992 "Image processing and its application to X-radiography." Conservation News No. 47, 37-39

Morgan N. 1991 An investigation into the plausibility of digitizing archaeological X-ray photographs and manipulating them using graphics imaging software in order that greater detail might be revealed. Unpublished BSc dissertation. Durham University

Gonzalez R.C. & Woods R.E. 1992 Digital Image Processing. New York. Addison-Wesley

APPENDIX 1

THE EQUIPMENT

This section is intended to provide specific suggestions for a basic image processing system for radiographs based on the IBM PC compatible computer. It should be born in mind that the choice of equipment is vast and the permutations endless therefore no direct comparison of equipment has been attempted. The system described was that appraised as most cost effective for the initial research work based on equipment and software readily available at the end of 1992.

Image acquisition

If a camera is being used to capture images from x-ray film then a good quality x-ray viewer is necessary. This should provide even background lighting and the means of masking out surrounding illumination.

A CCD camera is possibly the most versatile means of capturing an image with the added advantage of being small and robust. A CCD camera is obtainable in two versions one set up for the European video system, based on the 50Hz, 626 lines CCIR specifications with extensions to include the PAL and SECAM methods of colour coding. The second version is for the American and Japanese video system which is based on the 60Hz, 525 line RS-170 specification with extensions to include the NTSC colour system. The two are not compatible and therefore care should be taken in obtaining the correct version, in Britain this would normally be the CCIR system. Typical specifications and points to note for such a camera are as follows;

Monochrome or colour? Monochrome cameras are much cheaper than colour and if used with a monochrome frame grabber give slightly sharper images. However colour image processing whilst expensive at present does offer the possibility of far greater control over image manipulation. This is not of course applicable to processing radiographs.

A resolution of at least 640 x 480 pixels, this would normally be 768 x 512 pixels. Greater resolution than this, whilst desirable, is expensive at the present time. It is essential that the camera is compatible with the frame grabber board and therefore a higher resolution camera requires a higher resolution board which will also be more expensive.

A CCD camera requires a lens, this is not normally supplied with the camera. Obviously a good quality lens is essential and a zoom facility for low magnification work (x 6) is desirable. The majority of CCD cameras use what is termed a 'C' mount type of fitting. If microscope work is envisaged a microscope adapter will also be required.

The main manufactures of CCD cameras are Sony, JVC .

The frame grabber board

The function of frame grabber boards was described earlier. They are available in many forms from basic digitisers of monochrome signals with perhaps enough on board memory to store one image to true colour boards complete with extra processing power to perform specific operations such as histogram equalisation in real time. There is an obvious difference in price between these types of boards ranging from £1,000 to £12,000. It is difficult to specify a minimum requirement however a board that is capable of digitising the image and driving an image monitor, is able to store two images and has facilities for input/output look up tables should satisfy most needs. Obviously the board must be compatible with the host computer and capable of accepting the signal from the specified camera. If there is a possibility of future expansion of the system another factor to consider is the compatibility with additional dedicated processing boards capable of performing convolutions, Fast Fourier Transforms, geometric processing etc.

Manufacturers of frame grabber boards and related items include Data Translationand Synoptics.

The computer

In very broad terms IBM PC compatibles can be classified by the type of CPU (central processing unit) they are built around and the speed at which they operate - CPU's are usually defined by numbers e.g. 80286. Almost all CPU's are manufactured by Intel who are constantly developing newer more powerful units. Thus the 80286 was superseded by the 80386 and then the 80486 and now the 80586 (also called the Pentium). Other firms such as Cyrix have started producing CPU units and these are also known as 386's or 486's however they are not always equivalent to the Intel specifications. Further developments in computer technology are increasing the speed at which the machines are capable of operating. There is therefore a bewildering array of machines available with different specifications and a decision must be made as to which are the most suitable for the application.

Image processing packages themselves do not require large amounts of memory. The manipulation of images however does consume memory and the need for thousands of computations to be carried out during processing requires speed. In addition a number of the latest packages are written to run under Windows (a Graphical User Interface program) and this does require a certain level of memory to run efficiently.

Mathematical computations are most efficiently performed by a machine fitted with a maths co-processor. These are optional extras on 286 and 386 machines but are built into the full Intel 486DX unit (but not in the 486SX units). The temporary storage and processing of images requires random access memory (RAM) which can be added to computers in the form of expanded or extended memory. The minimum specifications for running Windows version 3.1 are considered to be a 386 machine with 4Mbytes of extended memory.

Considering the above requirements there are two options depending upon whether the chosen image processing package does or does not require Windows (see below for discussion of programs). If Windows is to be used the choice would appear to be a full 486 machine (i.e. with maths co-processor) running at 33MHz with at least 4Mbytes of extended memory although 8Mbytes should be considered preferable. If Windows is not to be used then either a 286 or a 386 machine could be used although they should be fitted with the appropriate maths co-processor. A further consideration for non Windows programs is that many are not able to access extended memory and some are restricted to the computers base memory.

Image processing packages do not consume large amounts of disk space therefore the size of the hard disk is dependant upon which other programs are to be run on that specific machine and how many images are to be stored. Most 486 machine are supplied as standard with around 100Mbtye of hard disk space and this should be adequate if the machine is to be used mainly for image processing although a 200Mbyte disk would allow for future development and/or storage of images whilst perhaps waiting for the optical disk market to standardise.

One final consideration is the size of the computer casing. The majority of frame grabber boards are full AT specifications and will not physically fit into many of the computer housing on sale at the present time.

Display

Display monitors have been discussed in section 2.2.4.

The program

The main advantages in using a program running under Windows is the ease of data transfer from one program to another. Applications in which this would be useful include the transfer from an image processing package of histogram data and grey level data to a spreadsheet for further analysis or the incorporation of images or LUT plots within a word processing package or presentation program.

THE IMAGE PROCESSING SYSTEM (DURHAM UNIVERSITY ARCHAEOLOGY DEPARTMENT)

The following equipment was used in the production of this report and related material.

Image acquisition All the radiographs were on Agfa D4 vacuum packed film and displayed on a RSL (Radiographic Supplies Ltd) Viewmaster X-ray viewer. A Sony DXC-151P CCD colour camera fitted with a Computar 18-108/2.5 zoom lens (up to 6 times magnification) was used for image acquisition.

Frame grabber board Data Translation DT2867-LC Precision Frame Grabber.

Computer Nektar 486 running at 33MHz with a 200Mbyte hard disk and 8Mbytes of RAM.

Image processing program Foster Findlay Associates PC_Image Plus Version 1.4 for Windows. The additional program modules and the demonstration program were produced with the aid of their C_Images Library routines.

Display monitor Sony Trinitron colour video monitor PVM-2130QM.

Hard copy production The photographs were produced using a Nikon FE SLR camera with Kodak ASA 100 colour print film. The monochrome laser prints were produced on a Hewlet Packard HP LaserJet 4M running at 600dpi except were stated.

APPENDIX 2

LOOK UP TABLES

The following LUTs are represented in graphical format in figures ?? to ?? and are included in a format suitable for the Foster Findlay PC_IMAGE programme on disk ? which accompanies this report. The LUTs are listed below with their figure name and number, their disk file name and comments on their effect.

Linear Greyscale - figure 6 **Disk filename** - LIN.LUT

Linear greyscale is the default LUT for maintaining the original order and distribution of the original grey values.

Inverse - figure 7 **Disk filename** - INVERT.LUT

Inverse LUT performs a linear remapping of the grey scale values such that 0 becomes 255 (black becomes white) and 255 becomes 0 (white becomes black).

Low grey level stretch #1 - figure 8 **Disk filename** - TEST1.LUT

Low grey level stretch LUT is designed to expand the original grey levels in the range 0 to 50 to new values between 0 and 110. This reveals details within the dark areas of the image.

Grey Mask - figure 9 **Disk filename** - GRYMASK1.LUT

Grey mask LUT maps the original grey scale values below 12 to 98. The remaining values remain as original. This produces a mid-grey background to the image.

Heated object - figure 10 **Disk filename** - HEAT.LUT

Heated object LUT represents the colour changes undergone by an object being heated through a wide temperature range. The scale runs from black to red (lowest grey values), through orange and yellow to white (the highest grey values).

High cut off - figure 11 **Disk filename** - HIGHCUT.LUT

The high cut off LUT is designed to highlight values above a grey level value of 200 in shades of blue/red. In addition the LUT expands the lower grey values from 0 to 30 in a similar way to that performed by the Low grey level stretch LUT (figure 8).

Sine wave - figure 12 **Disk filename** - SIN1.LUT

The sine wave LUT maps each colour component as a separate sine wave 6 degrees out of phase with each other. Altering the frequency and the phase of the sine waves will emphasis different grey scale ranges in different colours.

Edge highlight - figure 13 **Disk filename** - EDGE1.LUT

This LUT is included purely as an example of a method for highlighting different grey scale ranges in specific colours. The ranges would need to be changed to suit the specific image problem.

Magenta - figure 14 **Disk filename** - MAGENTA1.LUT

The magenta LUT is based on the fact that the human visual system can be most sensitive to hue changes within the magenta region.

Magenta with mask - figure 15 **Disk filename** - MAGENTA2.LUT

As above with the exception of mapping the 0 to 12 values to 98 producing a mid-grey background to the image.

Coloured map contour - figure 16 **Disk filename** - RAIBOW1.LUT

This colour scale uses a similar convention to colour contouring of maps. Blue through green is assigned to the lowest values, then through yellow and brown to mauve and white. The order of the original values is preserved whilst the changes are more easily recognised.

Coloured map with mask - figure 17 **Disk filename** - RAINBOW2.LUT

As above with the exception of mapping the 0 to 12 values to 98 producing a mid-grey background to the image.

Binary LUT removing background - figure 18 **Disk filename** - BIN1.LUT

The selected grey scale range is emphasised in white whilst the background is removed to a neutral grey.

Binary LUT preserving background - figure 19 **Disk filename** - BIN2.LUT

The selected grey scale range is emphasised in white whilst the background detail is preserved as in the original image.

APPENDIX 3

PROGRAM MODULES

The programme module ARCH_X is included on disk ? which accompanies this report. The programme is designed to run in conjunction with Foster Findlay's PC_IMAGE PLUS and includes the following additional processing functions;

Local Eq menu

7 to 1

This is a local equalisation function which operates over a 7 pixel by 7 pixel area writing the result to the central pixel.

11 to 1

As above with the exception of using a 11 pixel by 11pixel mask.

7 to 7

This is a local equalisation function which operates over a 7 pixel by 7 pixel area writing the results to a corresponding 7 by 7 area in the destination image. The source and the destination images must therefore be different or an error will occur.

11 to 11

This function operates as the 7 to 7 function with the exception of using an 11 pixel by 11 pixel area. Again the source and the destination image must be different or an error will occur.

Masks menu

Mask < 12 > 0

This function operates by assigning 0 grey level to all pixels in the range 0 to 12. It is designed to provide a constant black background to the image. The operation differs from using LUTs to remap the grey levels in that the original image is irretrievably altered.

Mask to grey 110

This function operates in the same manner as the previous one with the exception of assigning a grey level of 110 to the range 0 to 12. This provides a mid-grey background to the image.

Mask to grey 70

Similar to the above functions with the exception of assigning a grey level of 70 thereby producing a slightly darker background.

Image menu

Invert

This function inverts the grey scale values within the image assigning 0 to 255 and 255 to 0. As with the previous Mask menu items the operation is different to using LUTs as the original image is irretrievably altered.

Local Eq 7 x 7

This is the same as the 7 to 7 function in the Local Eq menu.

Av. + St. Dev

This function uses a 7 pixel by 7 pixel mask from which the average grey level and the standard deviation are calculated for that particular image area. The result is assigned to the central pixel of the area.

To use ARCH_X both PC_IMAGE and ARCH_X must be run simultaneously. In PC_IMAGE load the required image to be processed, from the OPTION menu use copy link DDE to put the image and ROI definitions into the Windows clipboard. In ARCH_X use the EDIT menu to paste these into that programme. The image processing functions in ARCH_X can now be used to alter the loaded image. When the processing is completed return to PC_IMAGE to either save the results or apply further functions.

APPENDIX 4

DISK IMAGES AND FILE NAMES

The following are the images, their respective file names and where appropriate the corresponding text figure number included on disk ?

Figure No.	Description	Filename	Comments
Fig 31	Gold pendant	Mkfull.imf	
Fig 36	Two fragments of iron	Drys1.imf	
Fig 38	As above after expansion of grey scale	Drys2.imf	
	Included in the demo programme only	Lut1.imf	
Fig 26 & 53	Copper alloy coin	Coin.imf	
Fig 27	Coin after histogram equalisation	Coineq.imf	
	Included in demo programme only	Coins2.imf	
Fig 57	Coin in fig 26 with surface texture effect	Cointex.imf	
Fig 54	Coin after high pass FFT	Coinfft,imf	
Fig 35	Four views of coin with different processing applied	Coin4.imf	
	Comparison of sword enhancement	Swcomp1.imf	In demo programme only
Fig 46	Roman iron padlock	Lock1.imf	
	Padlock after filtering	Lock1add.imf	In demo programme only
	Padlock after cleaning	Lockrl.imf	In demo programme only
	Text file overlay included in demo programme	Title2.bin	Binary text file for overlay
	"	Drys1.bin	"
	"	Lut1.bin	"
	"	Dry2.bin	"
	"	Cntx1.bin	"
	"	Cnftx1.bin	"
	"	Coin4tx.bin	"
	"	Swcomp1.bin	"
	"	Lock1add.bin	"
	"	Lckadd.bin	"

APPENDIX 5

THE DEMONSTRATION PROGRAM

The demonstration programme is called DURDEM1.EXE and is designed to run under Windows. The programme and related images files must be copied into a directory named 'C:\IMAGES' . To operate load windows and run FILE MANAGER. From the FILE menu choose the RUN option and type 'C:\DURDEM1'. Once the programme has loaded choose 'READ DISK IMAGE' from the IMAGE menu; a message box will appear from which the demo can be started. The programme will display a series of annotated images illustrating a number image processing techniques as applied to archaeological radiographs.

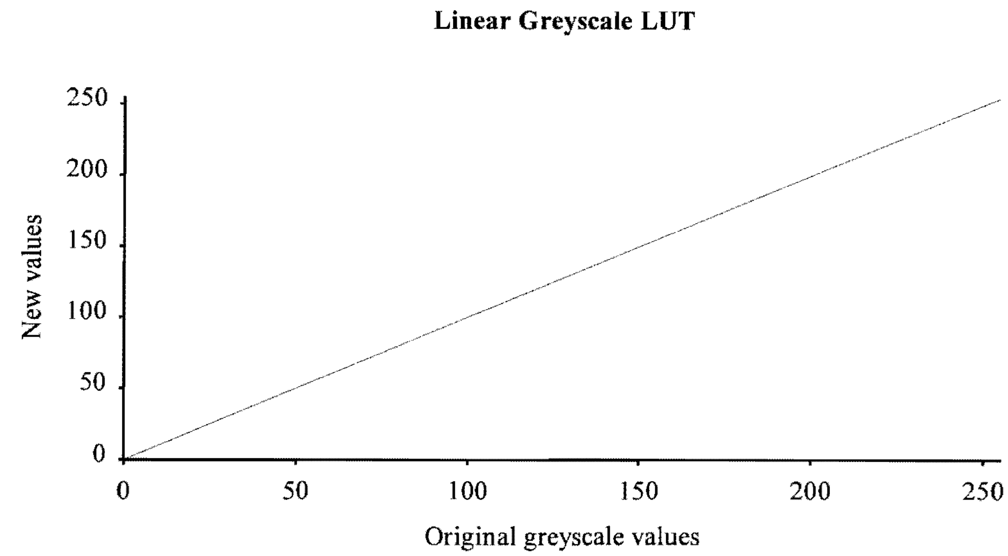

Figure 6 *Linear greyscale is the default LUT for maintaining the original order and distribution of the original grey values.*

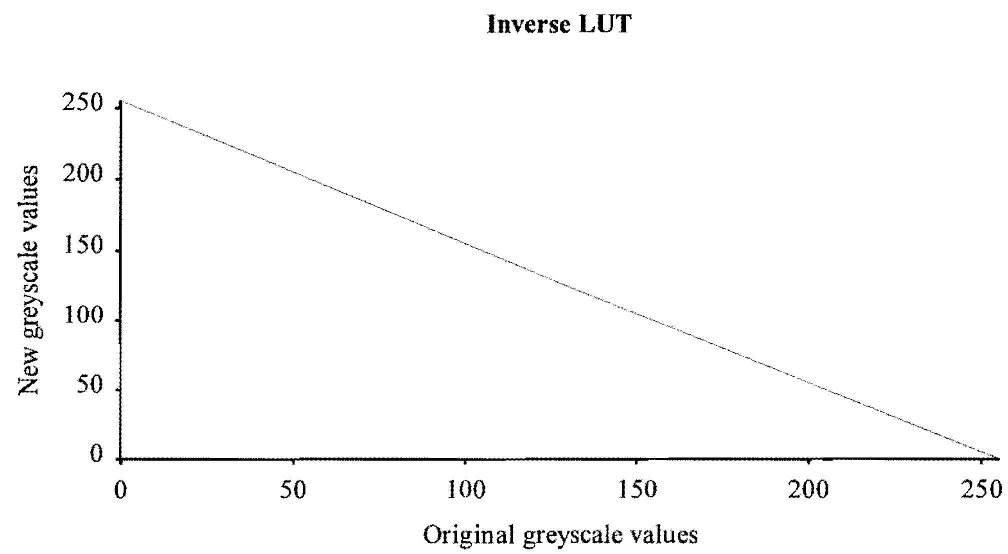

Figure 7 *Inverse LUT performs a linear remapping of the grey scale values such that 0 becomes 255 (black becomes white) and 255 becomes 0 (white becomes black).*

Low Grey Level Stretch #1

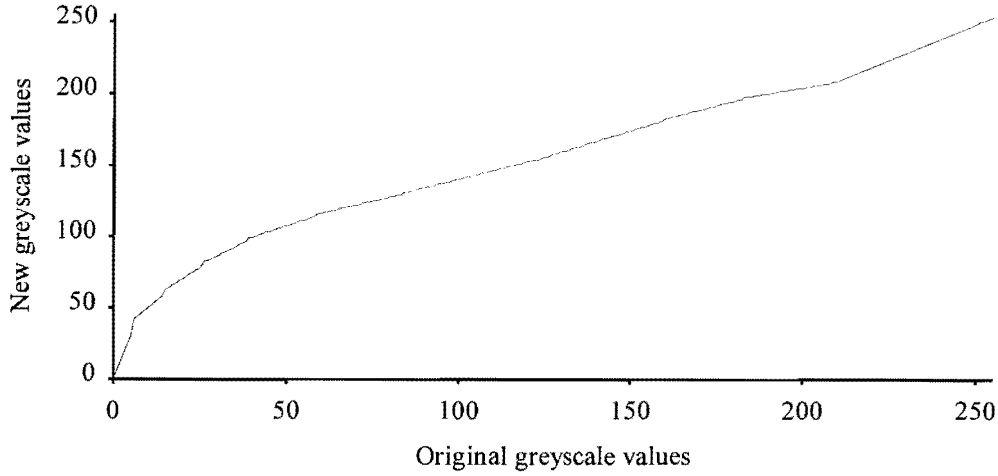

Figure 8 *Low grey level stretch LUT is designed to expand the original grey levels in the range 0 to 50 to new values between 0 and 110. This reveals details within the dark areas of the image.*

Grey Mask LUT

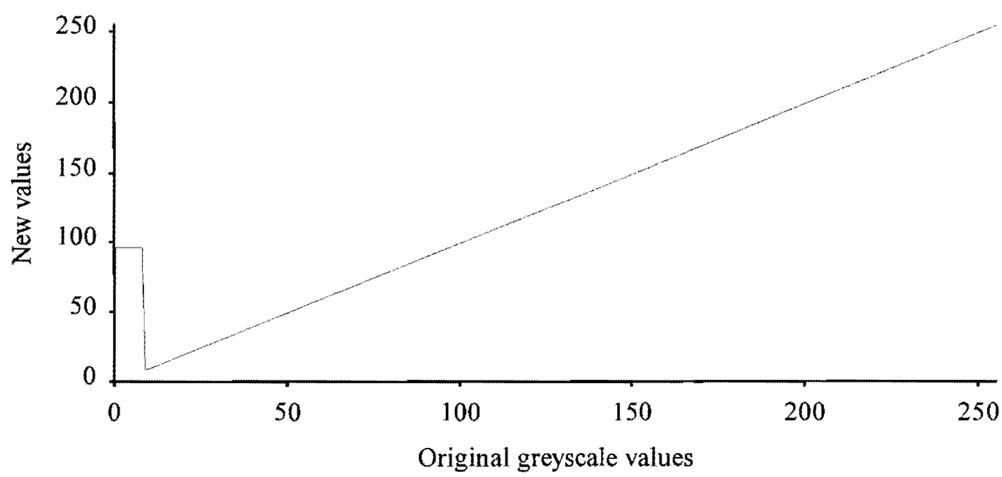

Figure 9 *Grey mask LUT maps the original grey scale values below 12 to 98. The remaining values remain as original. This produces a mid-grey background to the image.*

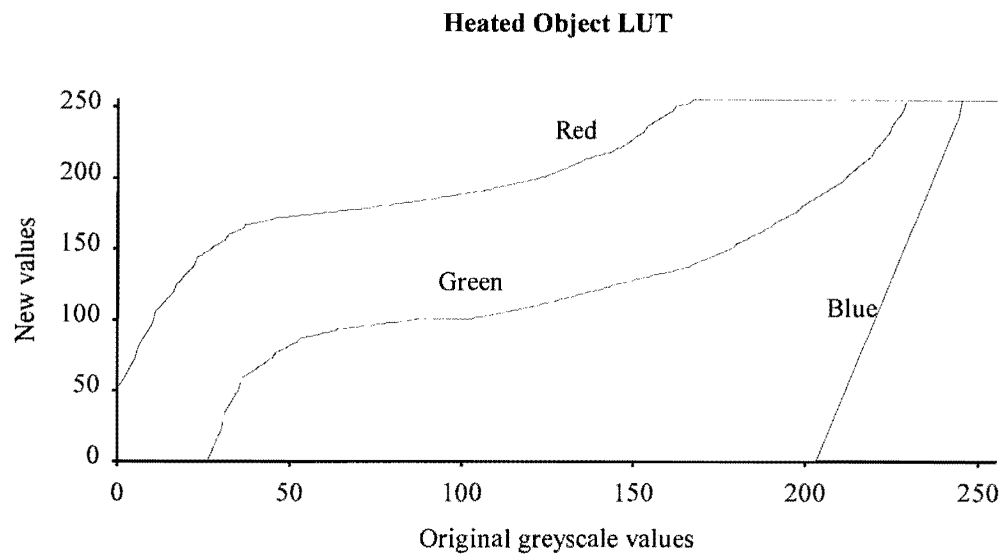

Figure 10 *Heated object LUT represents the colour changes undergone by an object being heated through a wide temperature range. The scale runs from black to red (lowest grey values), through orange and yellow to white (the highest grey values).*

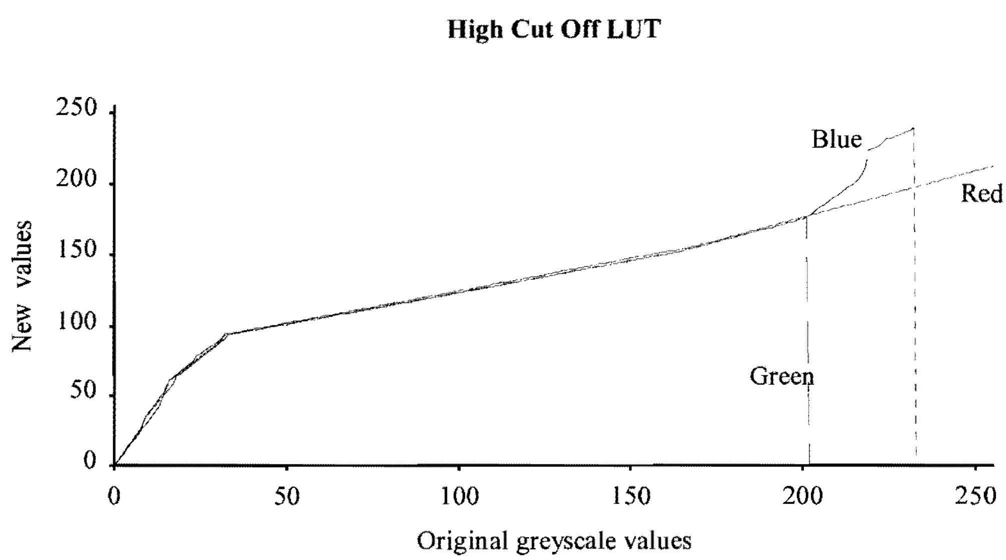

Figure 11 *The high cut off LUT is designed to highlight values above a grey level value of 200 in shades of blue/red. In addition the LUT expands the lower greyvalues from 0 to 30 in a similar way to that performed by the Low grey level stretch LUT (figure 8).*

Sine Wave LUT

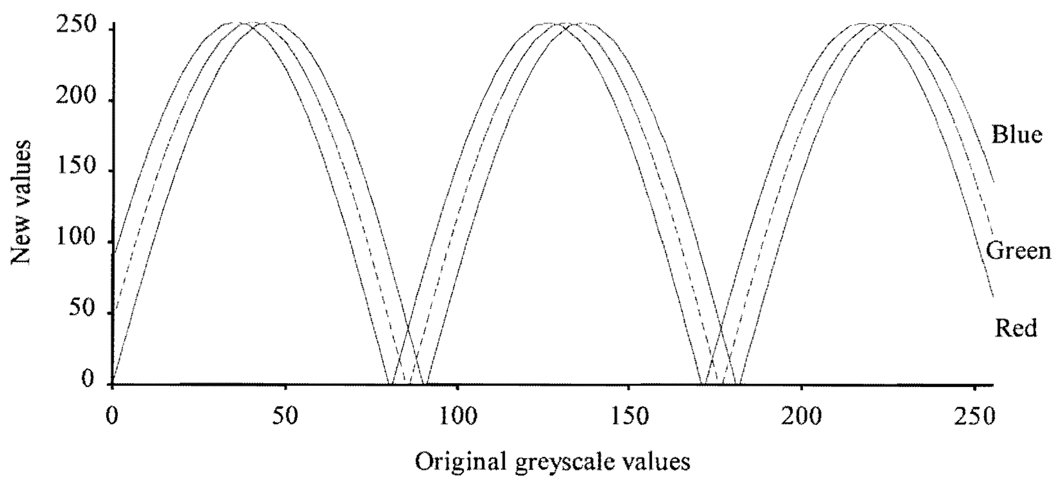

Figure 12 *The sine wave LUT maps each colour component as a seperate sine wave 6 degrees out of phase with each other. Altering the frequency and the phase of the sine waves will emphasis different grey scale ranges in different colours.*

Edge Highlight LUT

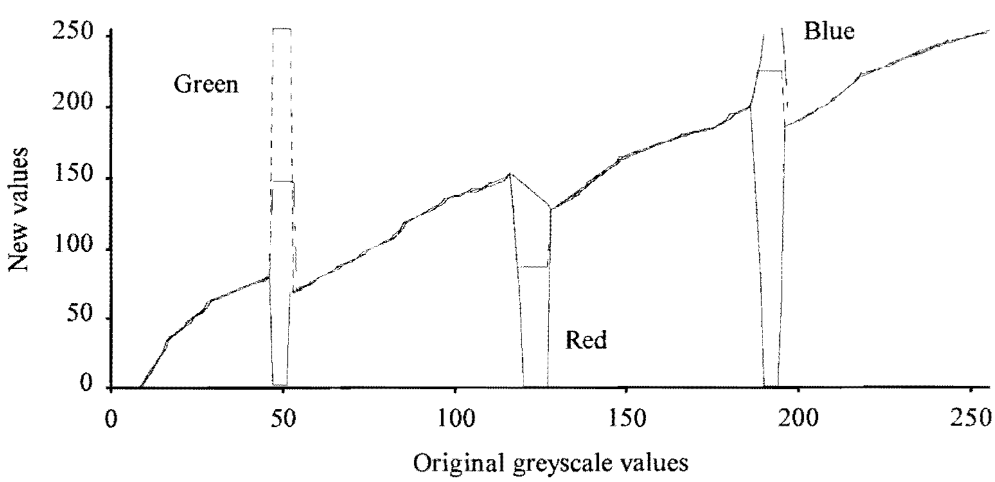

Figure 13 *This LUT is included purely as an example of a method for highlighting different grey scale ranges in specific colours. The ranges would need to be changed to suit the specific image problem.*

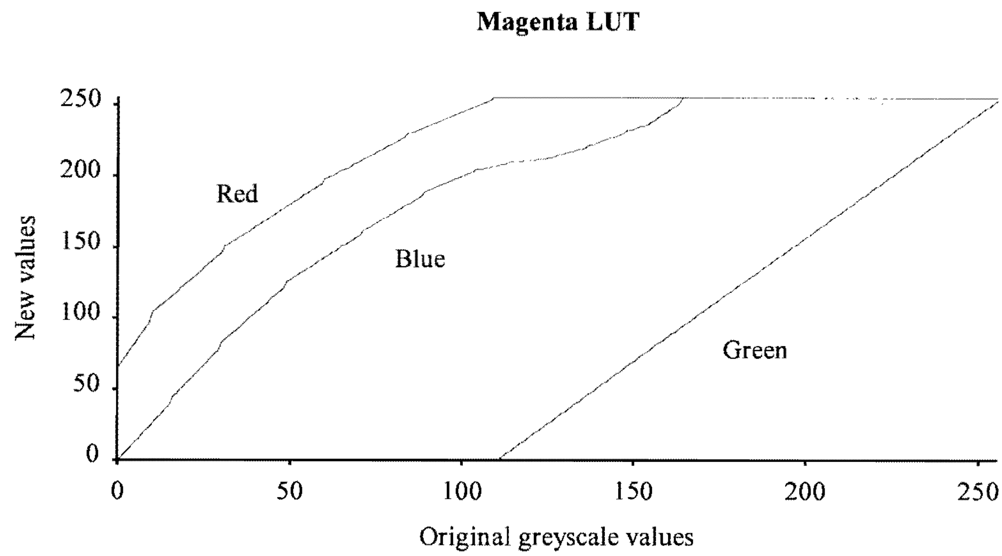

Figure 14 *The magenta LUT is based on the fact that the human visual system can be most sensitive to hue changes within the magenta region.*

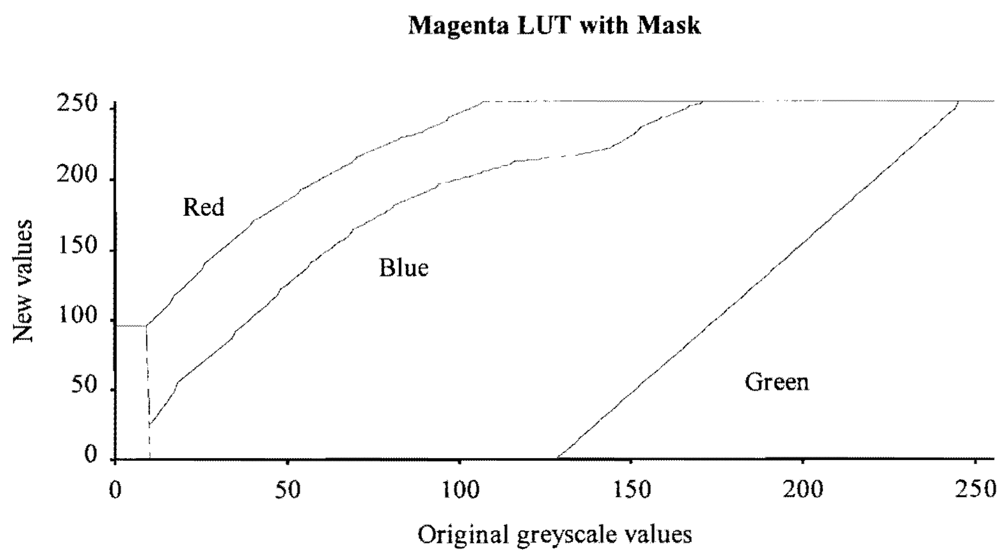

Figure 15 *As above with the exception of mapping the 0 to 12 values to 98 producing a mid-grey background to the image.*

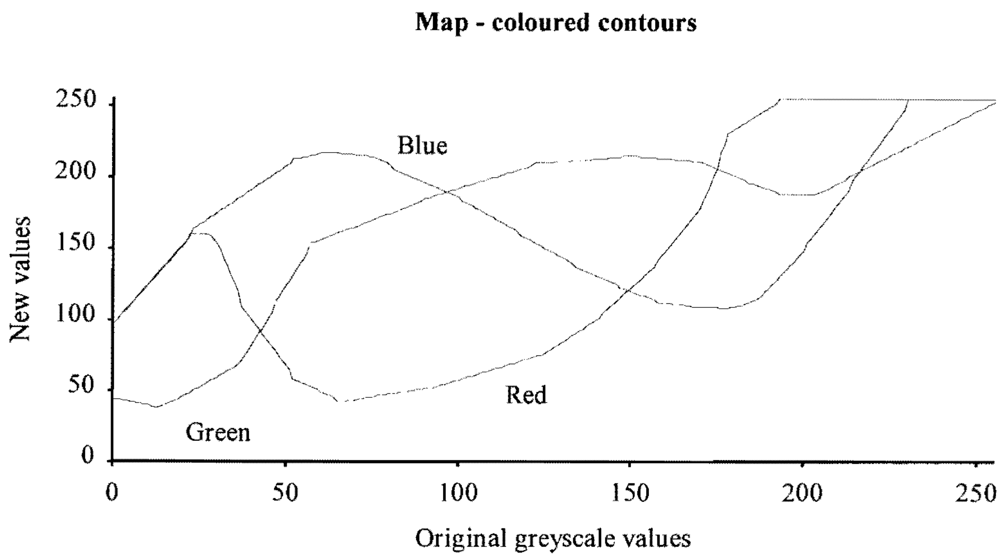

Figure 16 *This colour scale uses a similar convention to colour contouring of maps. Blue through green is assigned to the lowest values, then through yellow and brown to mauve and white. The order of the original values is preserved whilst the changes are more easily recognised.*

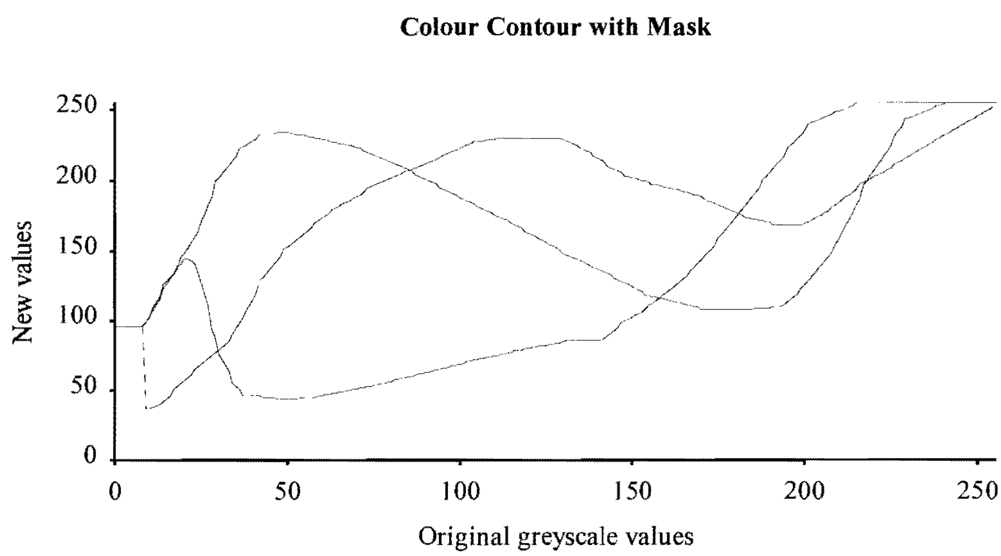

Figure 17 *As above with the exception of mapping the 0 to 12 values to 98 producing a mid-grey background to the image.*

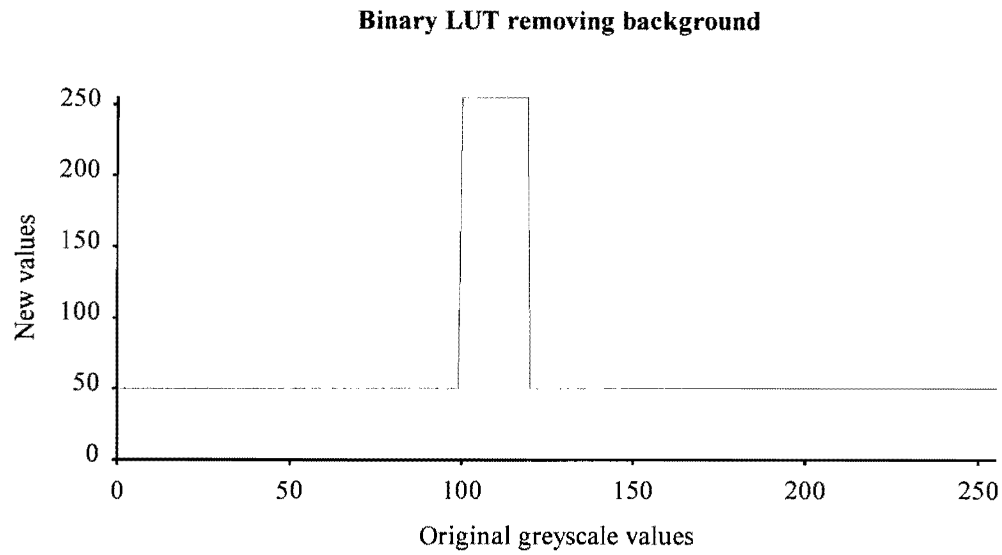

Figur 18 *The selected grey scale range is emphasised in white whilst the backround is removed to a neutral grey.*

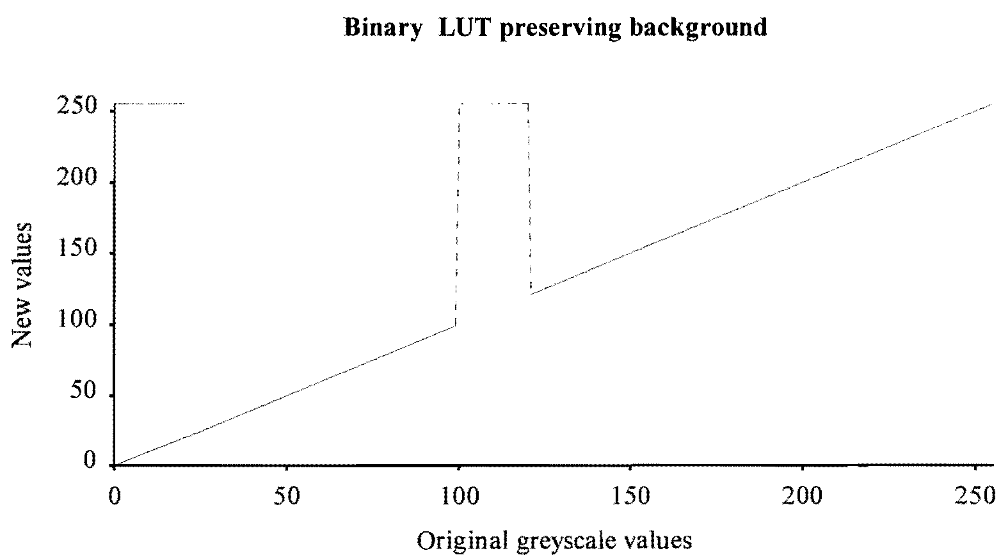

Figure 19 *The selected grey scale range is emphasised in white whilst the background detail is preserved as in the original image.*

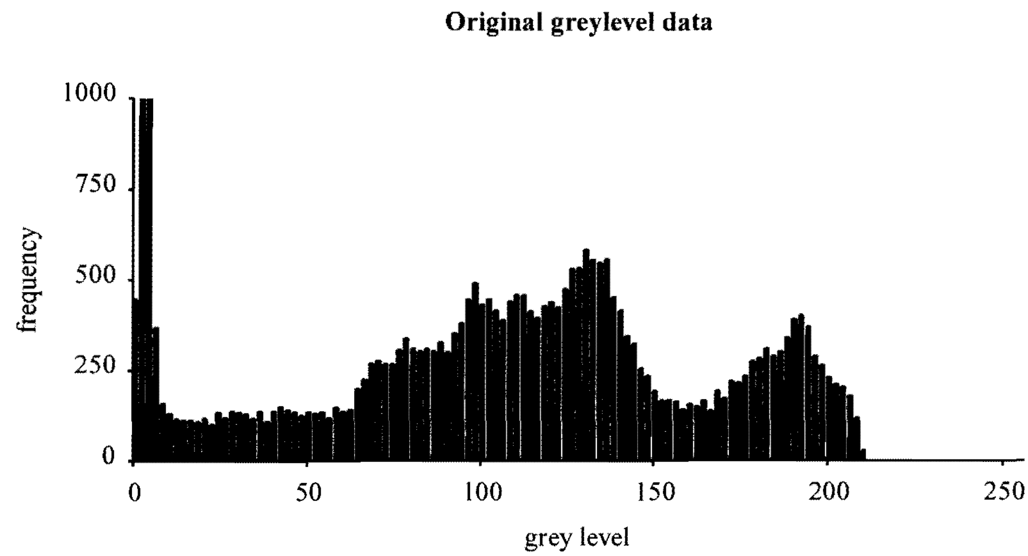

Figure 20 *Histogram of the original grey level data from the coin.*

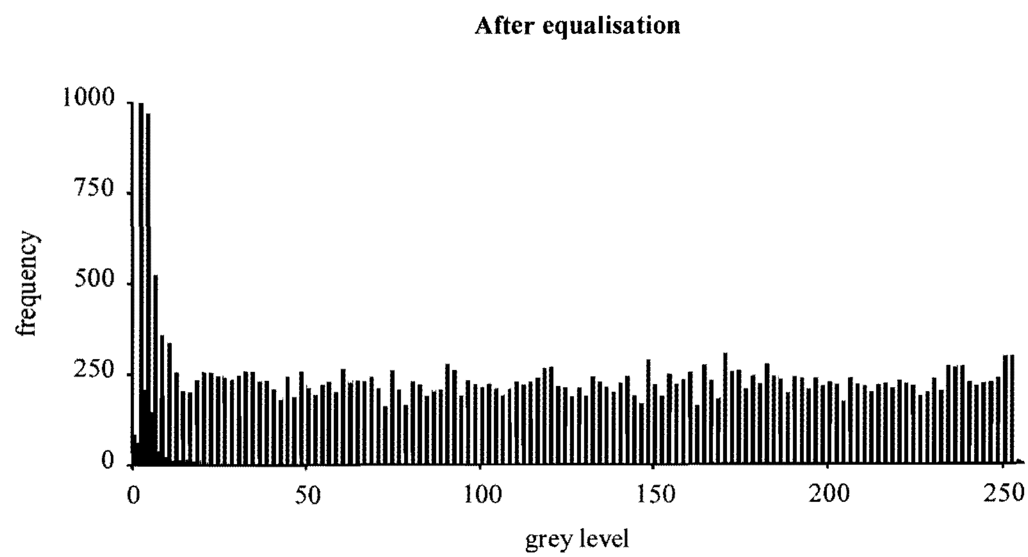

Figure 21 *A histogram of the grey levels from the coin after histogram equalisation.*

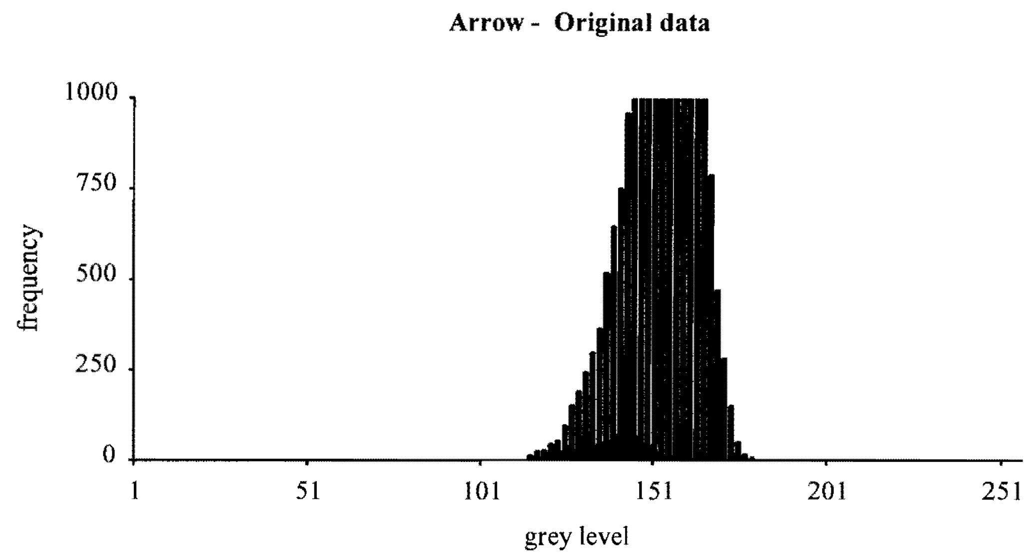

Figure 22 *Original histogram of the image grey levels.*

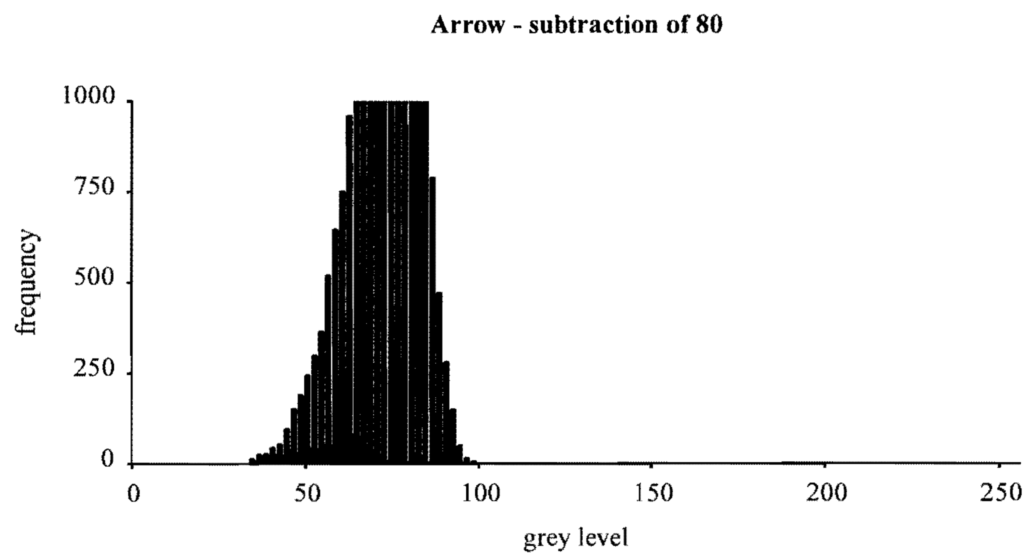

Figure 23 *The effect of histogram slide - a value of 80 has been subtracted from all grey level values.*

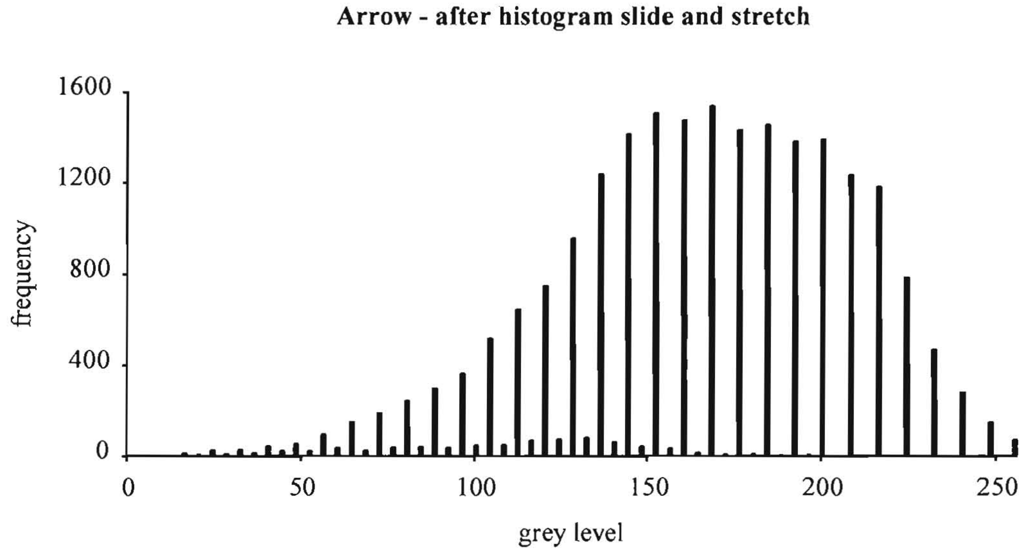

Figure 24 *Histogram produced by using the histogram slide and stretch technique. A value of 80 was subtracted from each of the original grey scale values and the resulting values multiplied by 2.*

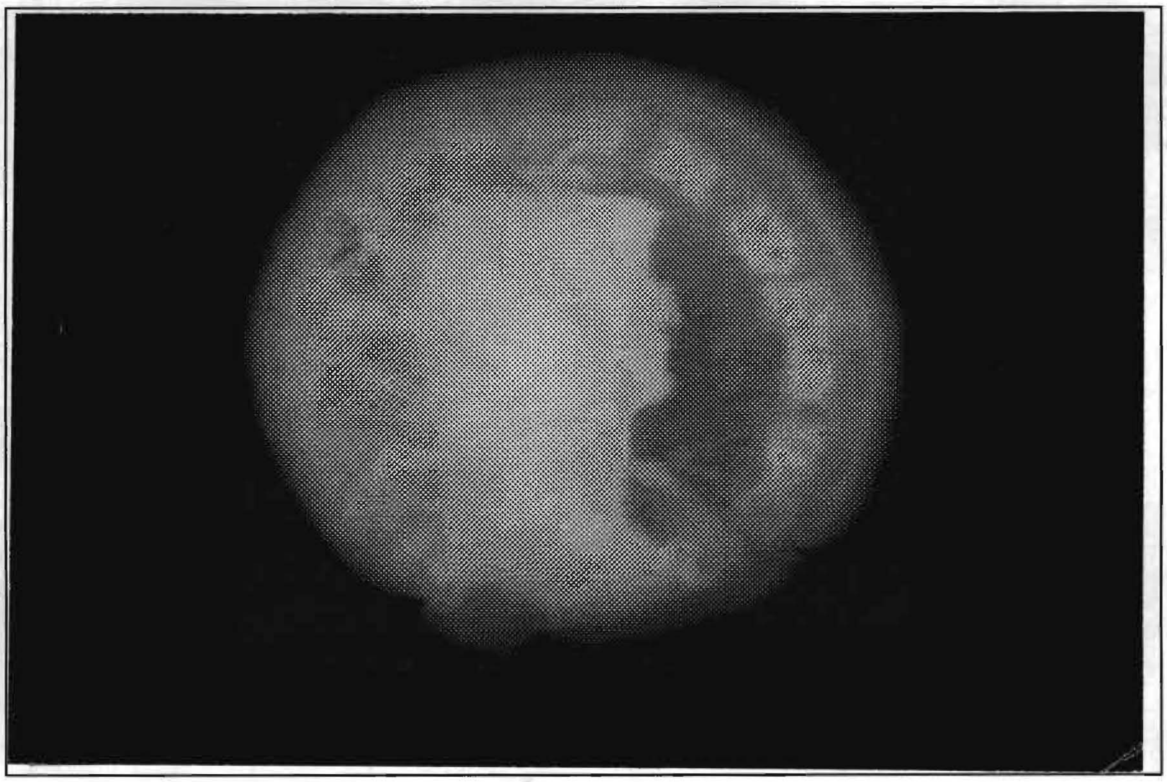

Figure 25 *Radiograph of coin printed at 300 dpi on a HP laserjet 3.*

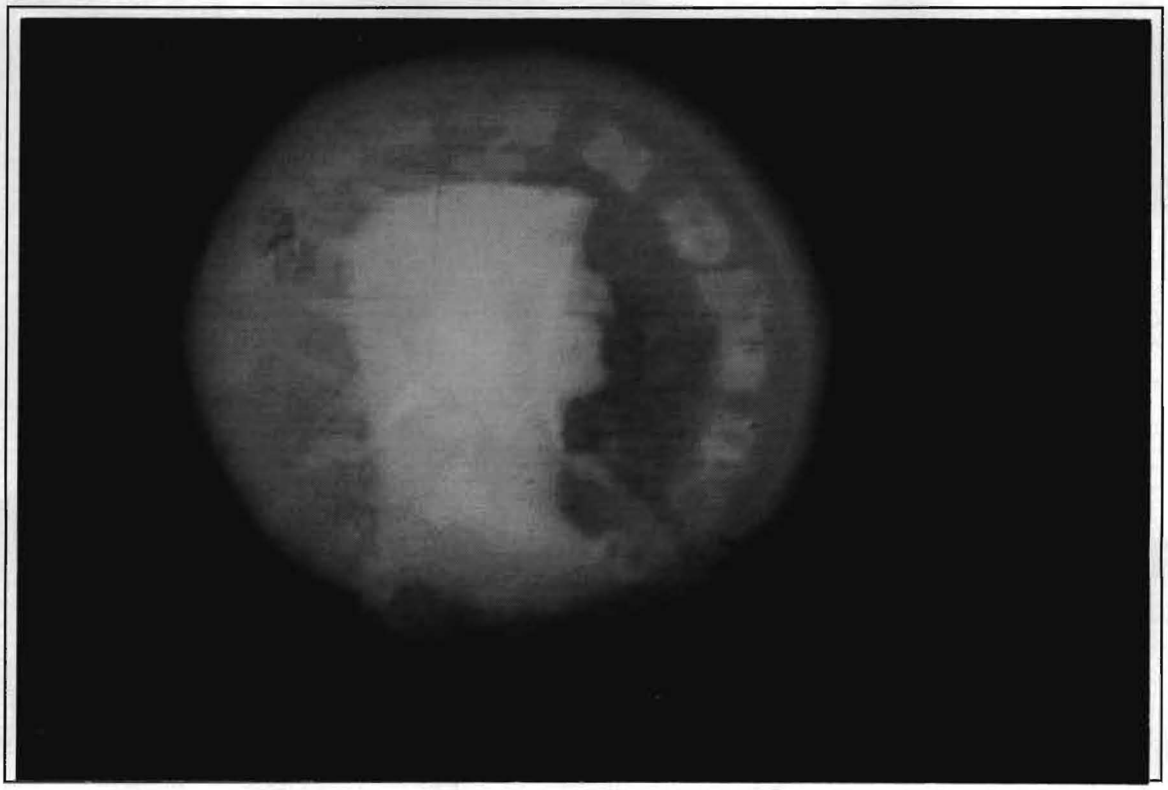

Figure 26 *Radiograph of copper alloy coin as figure 25 printed at 600 dpi on a HP laserjet 4M*

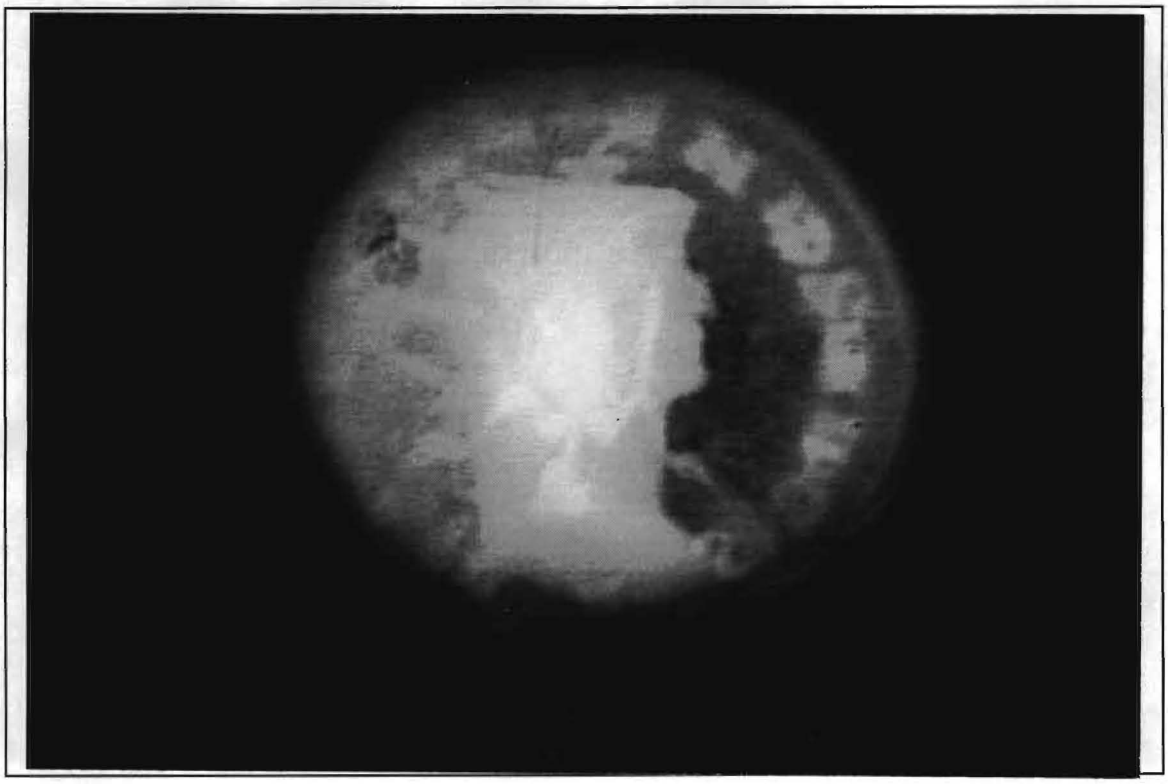

Figure 27 *The coin in figure 26 after histogram equalisation. The detail is more visible.*

Figure 28 *Original radiograph of a panel painting.*

Figure 29 *Figure 28 after histogram equalisation.*

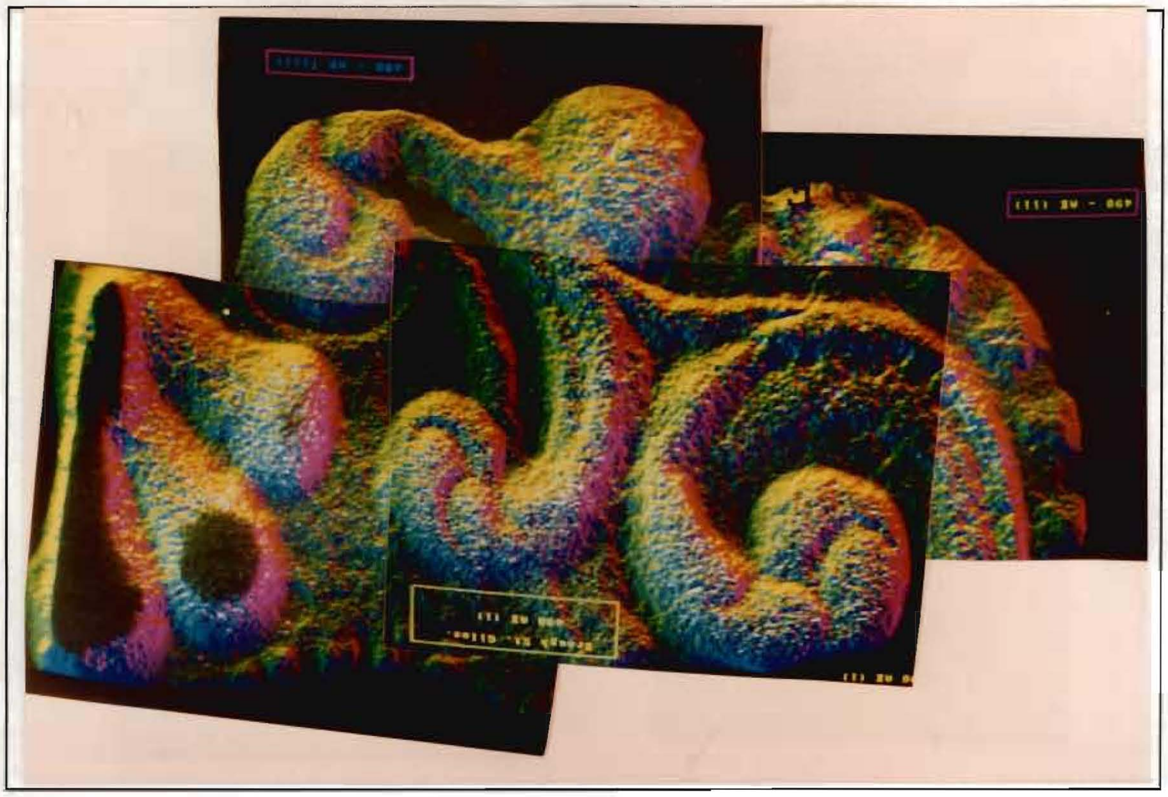

Figure 30 *Enhanced radiograph with false colours giving a 3D representation of the artefact.*

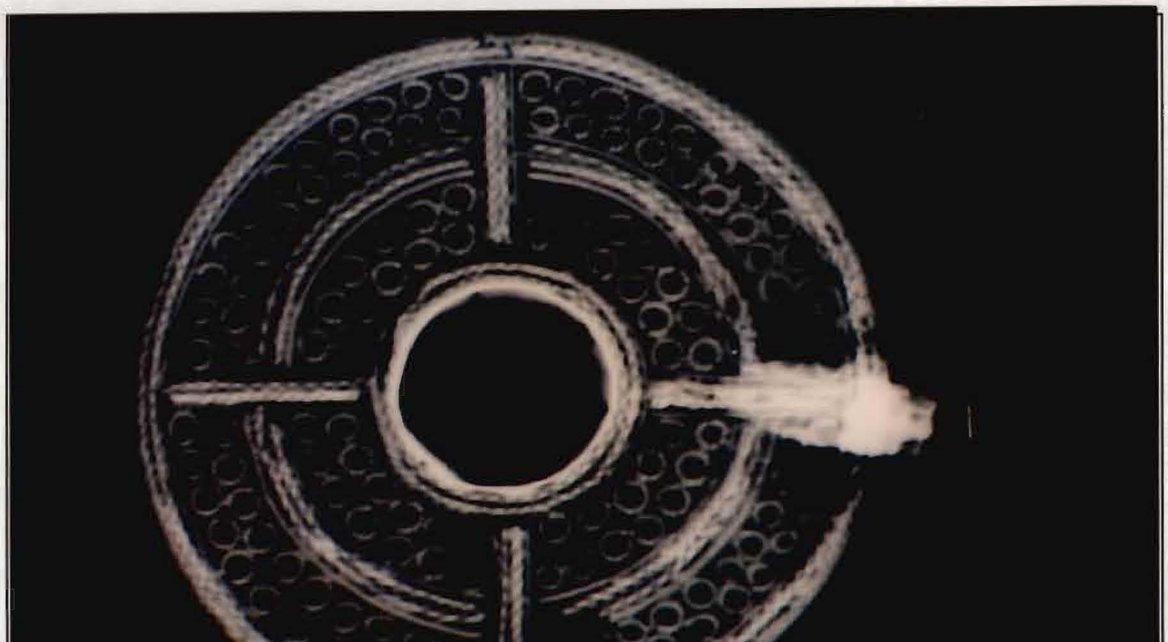

Figure 31 *Original radiograph of a gold pendant measuring 2.5 cm in diameter.*

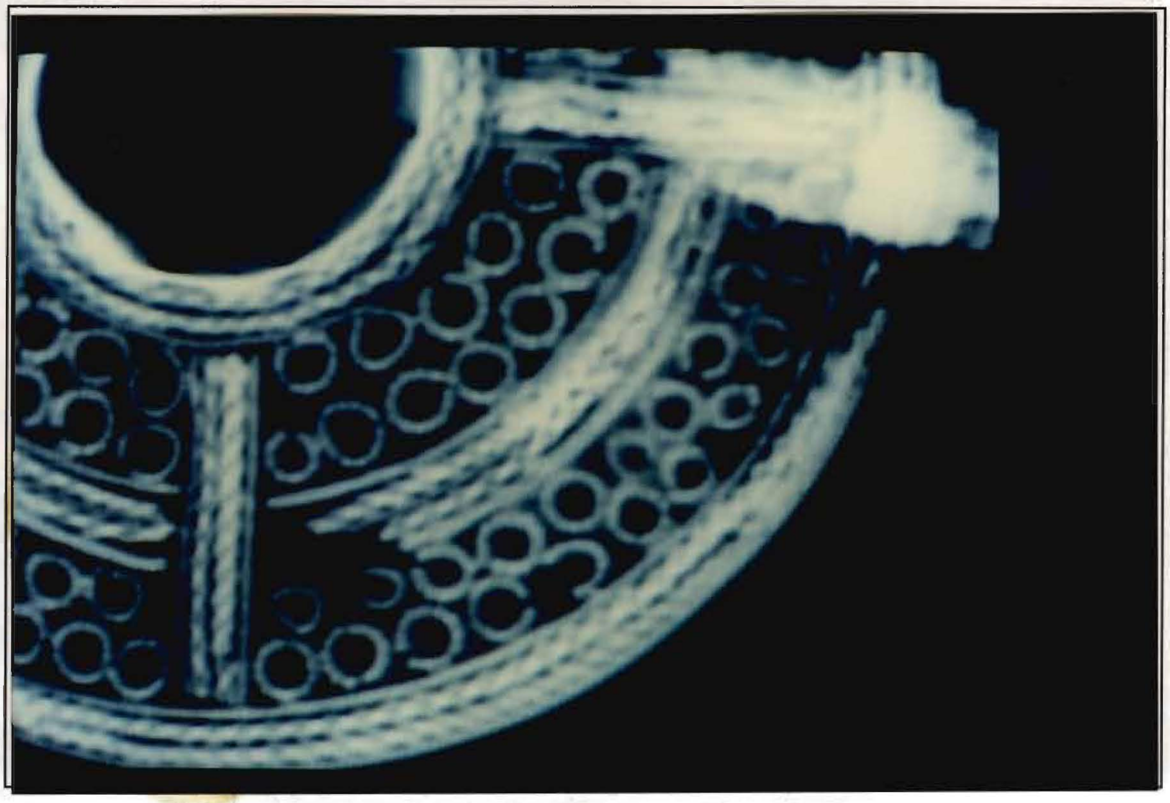

Figure 32 *A section of the pendant after 2 x magnification of the digital image*

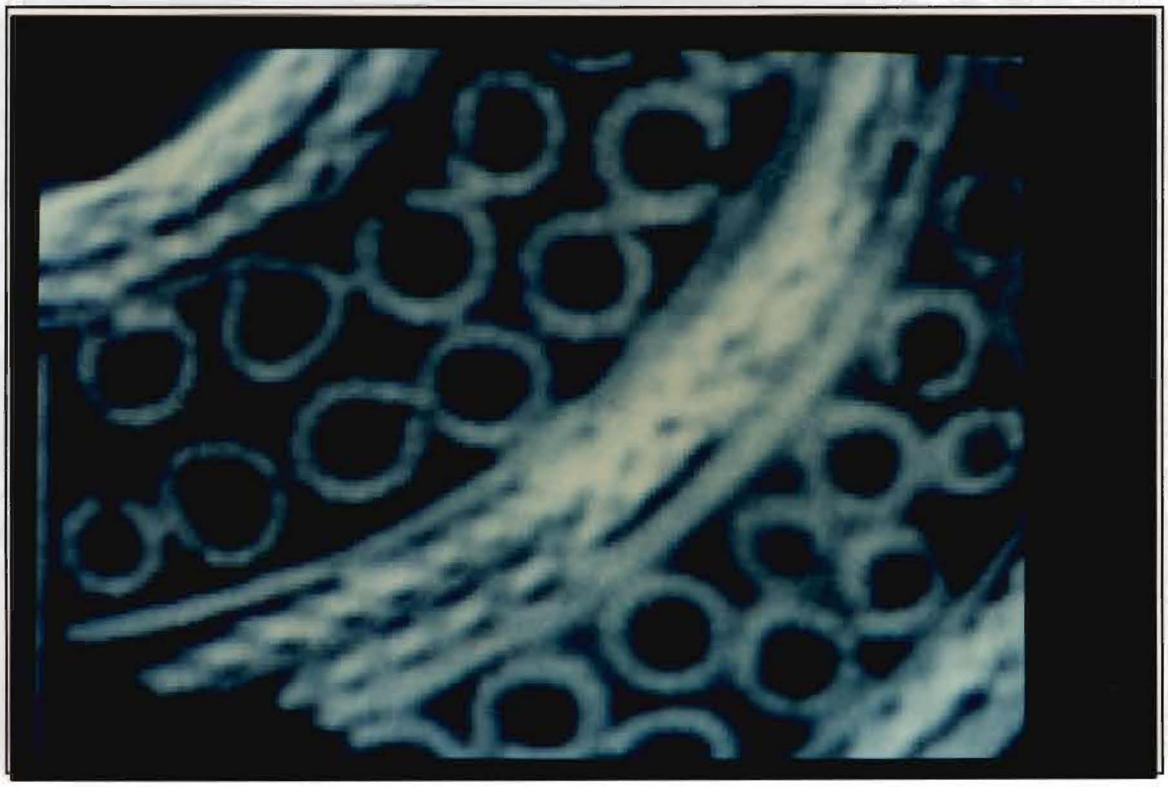

Figure 33 *The same section as in figure 32 with 4 X magnification. The individual pixels are quite clearly visible.*

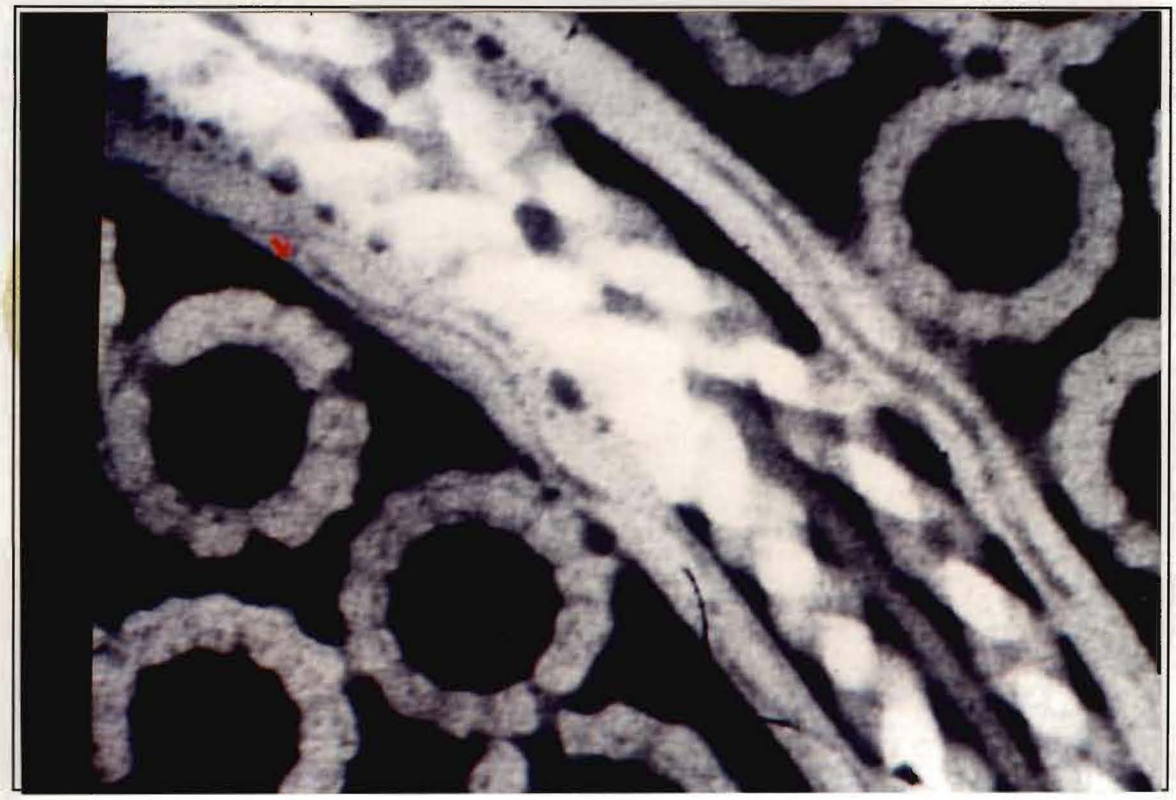

Figure 34 *A section of the original radiograph digitised from a magnified image using a binocular microscope.*

Figure 35 *Four images reduced to 25% of their area. Useful for comparing images and in presentation.*

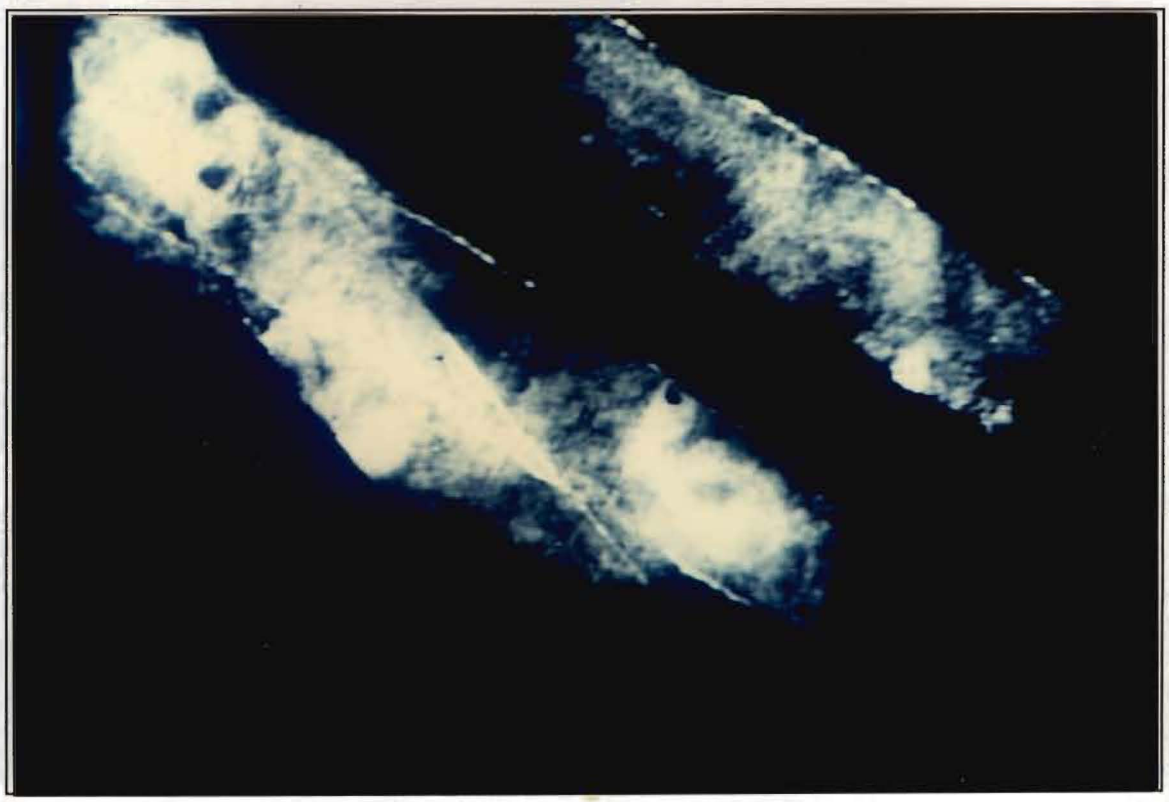

Figure 36 *Original radiograph of two fragments of iron. Approx. 4 times original size.*

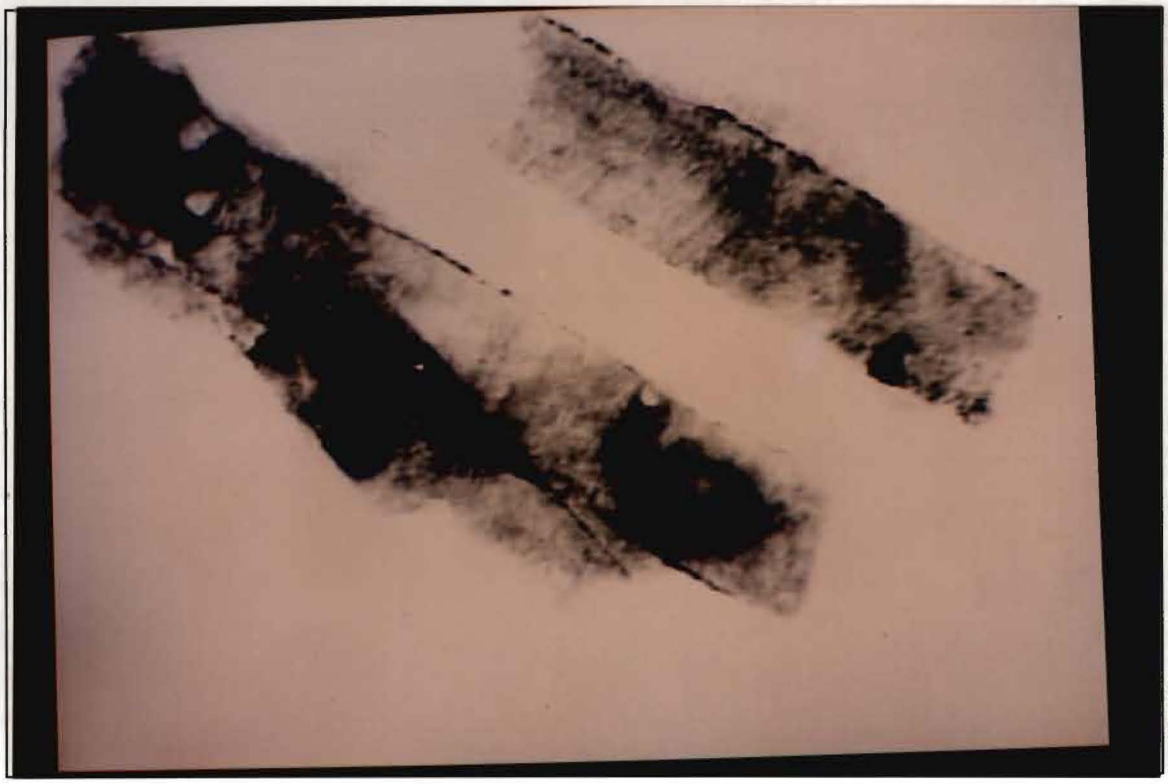

Figure 37 *Inversion of the grey scale in figure 36 using the LUT illustrated in figure 7.*

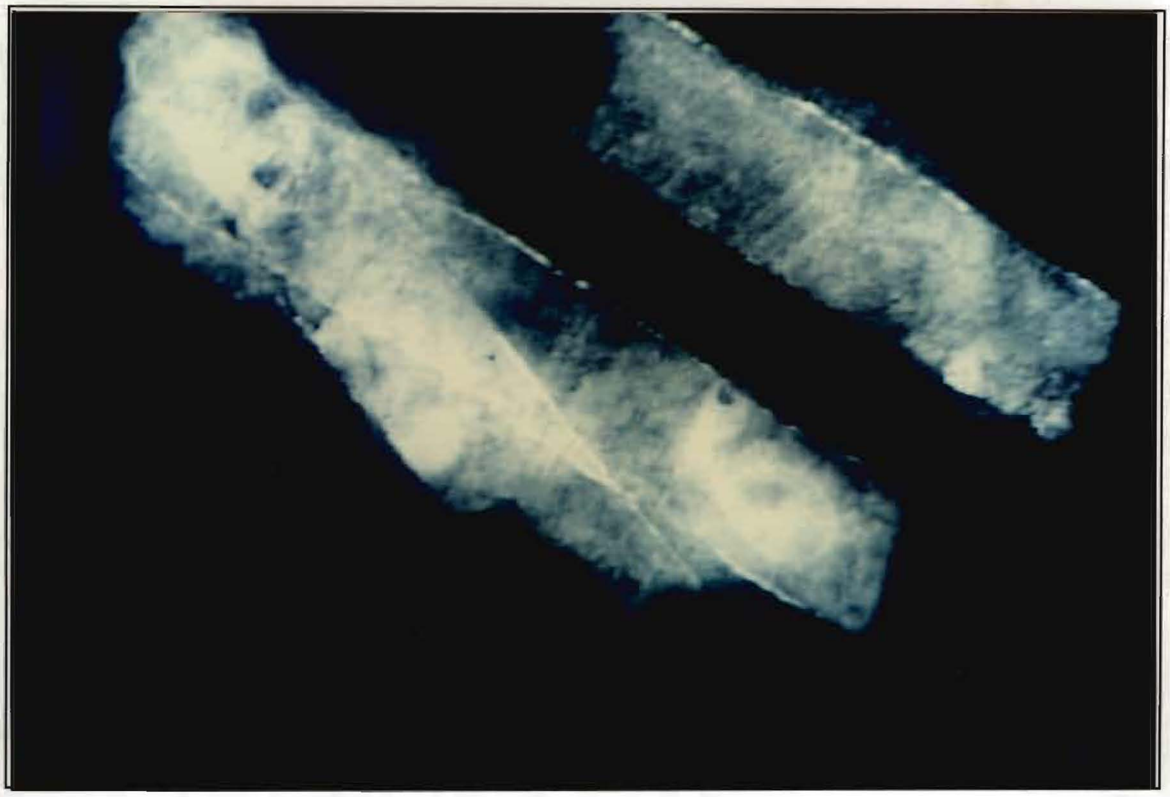

Figure 38 *Expansion of the grey scale using LUT in figure 8 provides more information in darker areas.*

Figure 39 *Original radiograph of iron arrowhead. The background has been masked to grey using the LUT in figure 9. The histogram of the grey levels is illustrated in figure 22.*

Figure 40 *The central section of the radiograph has under gone a histogram slide and stretch revealing more detail than in the original. Figure 24 shows the new histogram.*

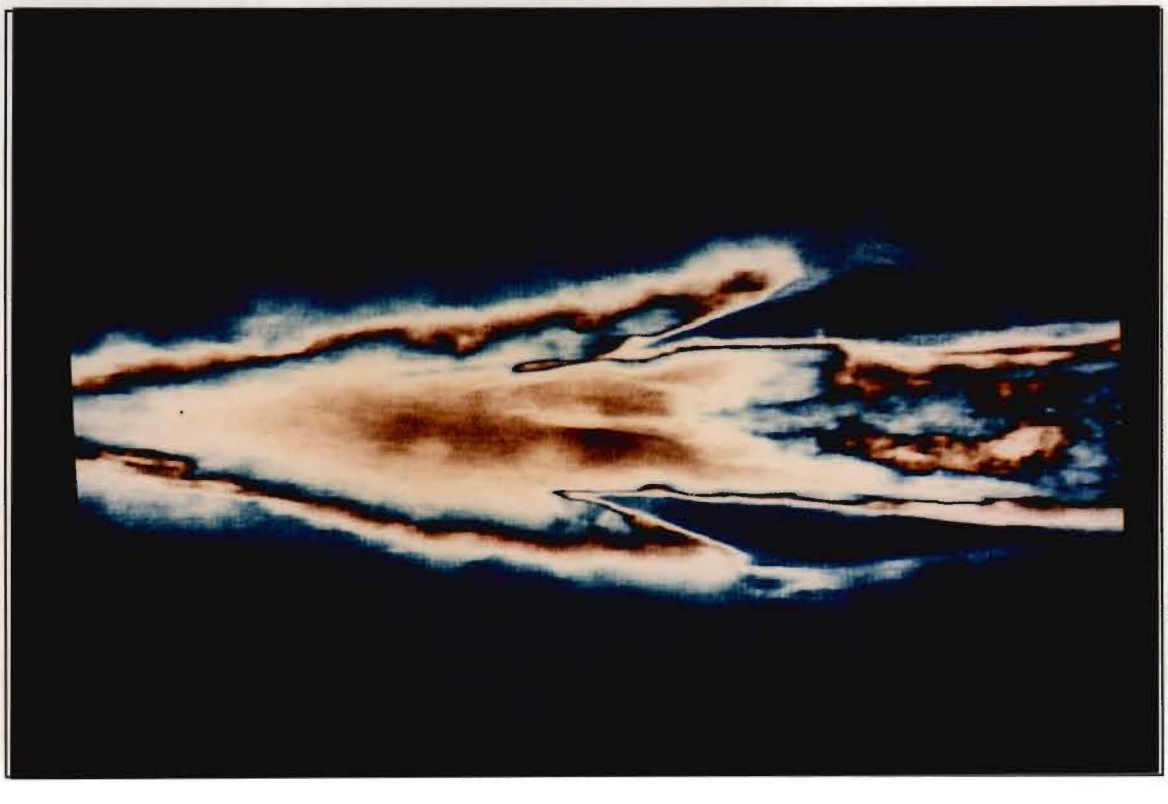

Figure 41 *A sine wave LUT (figure 12) has been used to highlight various areas .*

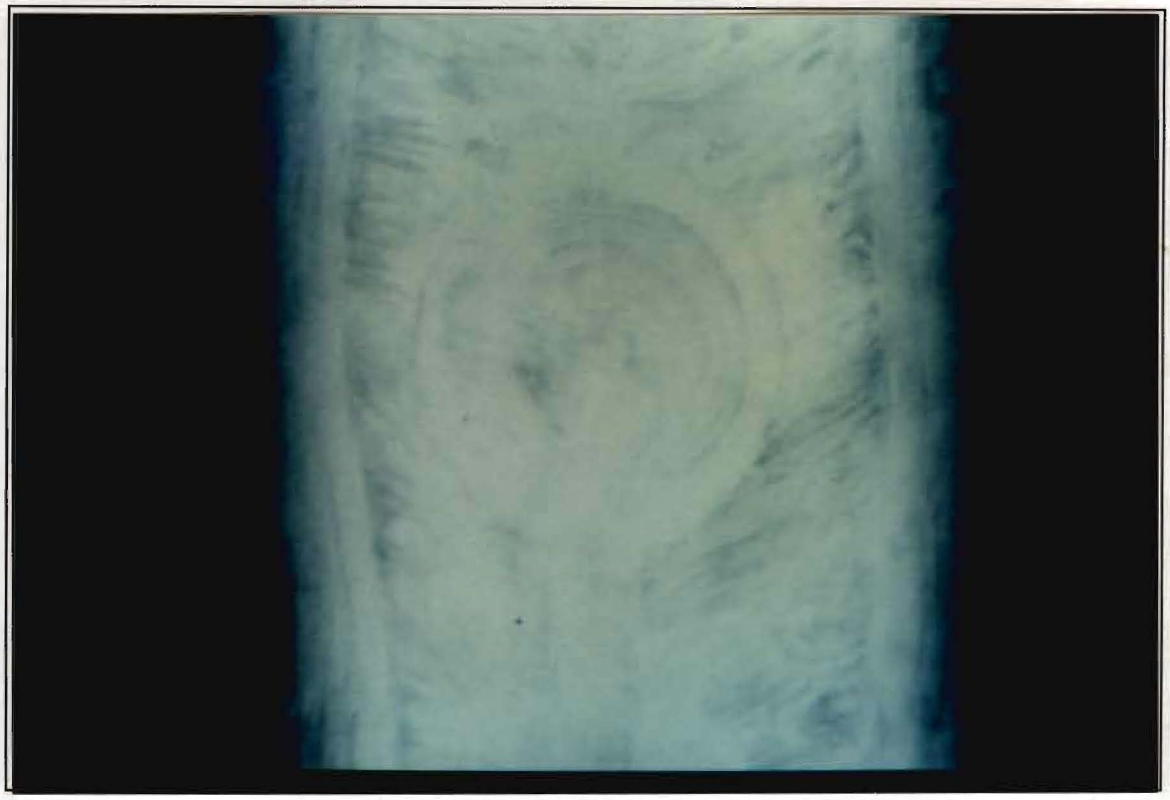

Figure 42 *Original radiograph of an iron sword blade showing the pattern welding and decorative inlay.*

Figure 43 *The radiograph of the sword after a 7 x 7 local equalisation.*

Figure 44 *A sharper image is produced using FFT to filter out the low frequency components.*

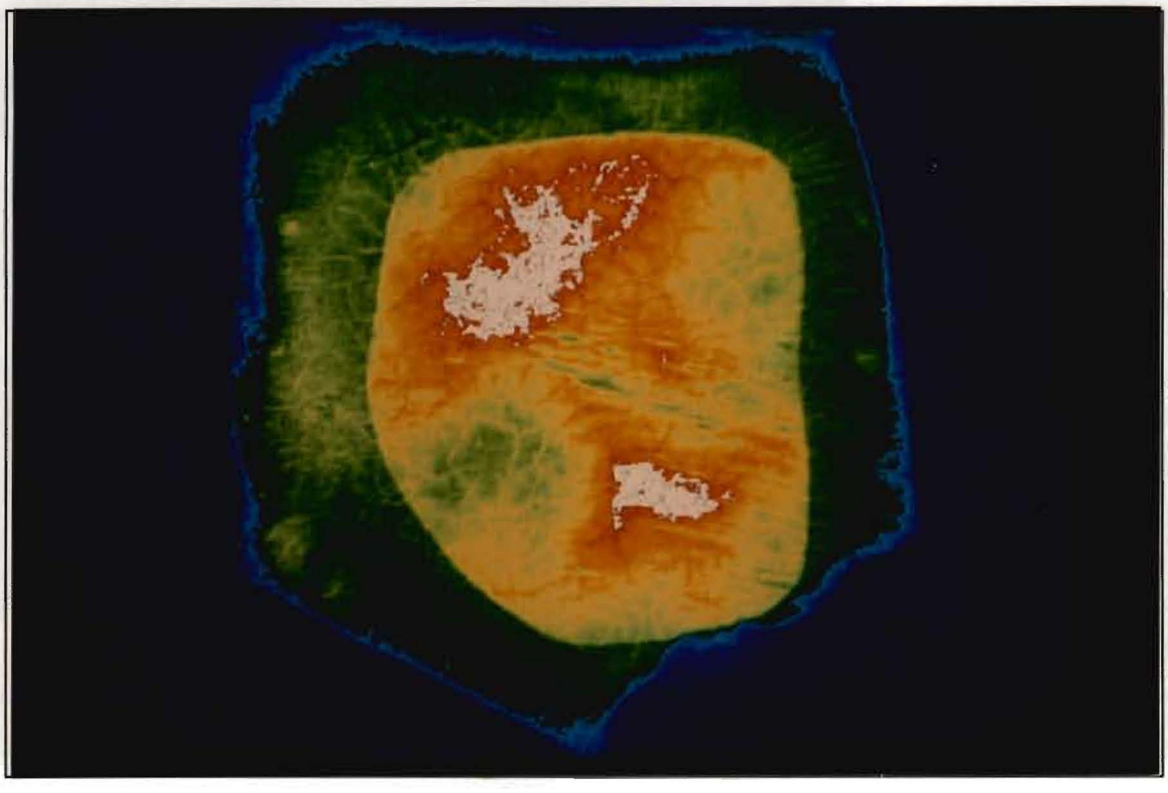

Figure 45 *Radiograph of a fragment of window glass enhanced with the application of the contour map LUT shown in figure 16.*

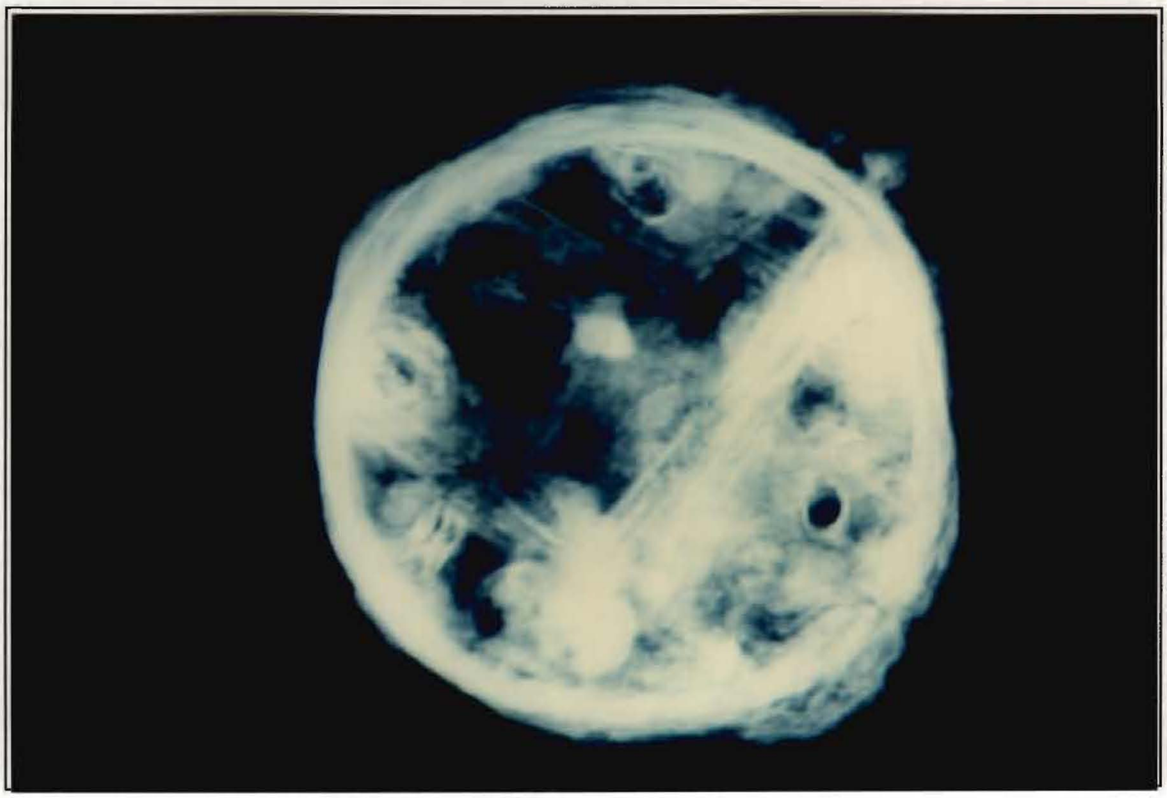

Figure 46 *Original radiograph of a iron padlock.*

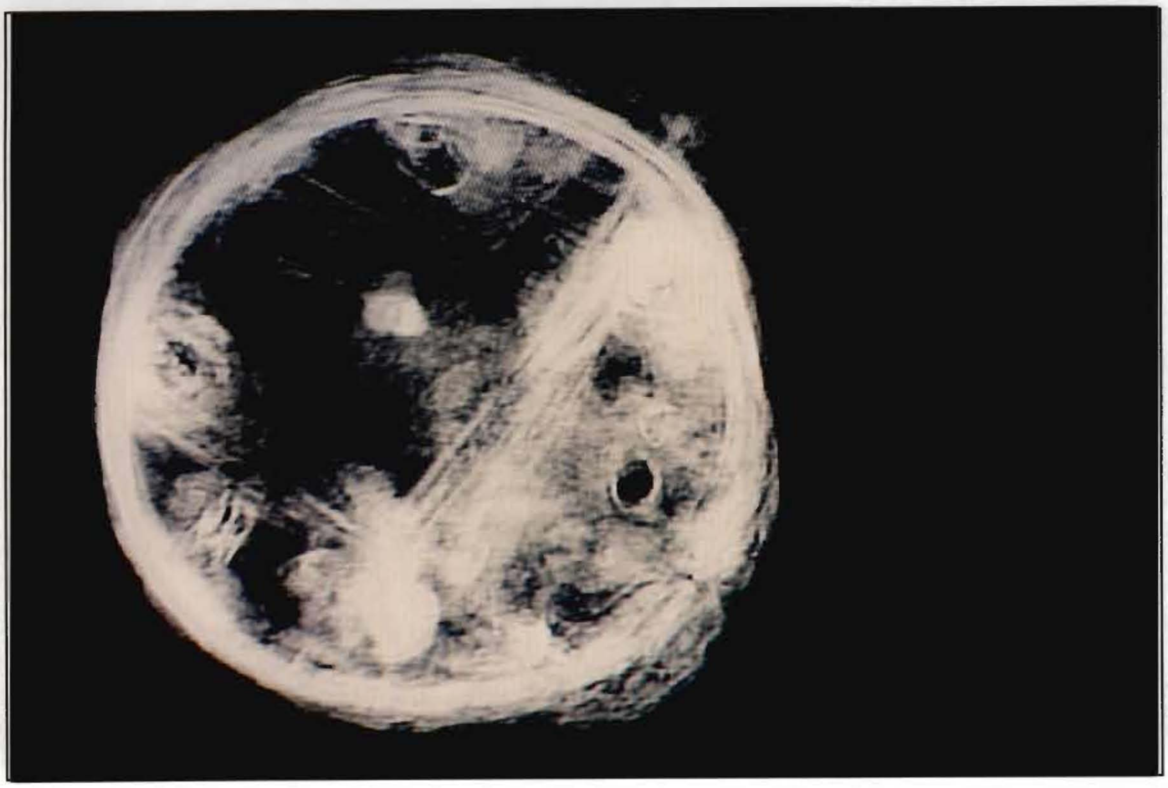

Figure 47 *The radiograph after application of the sharpening filter illustrated in figure 1.*

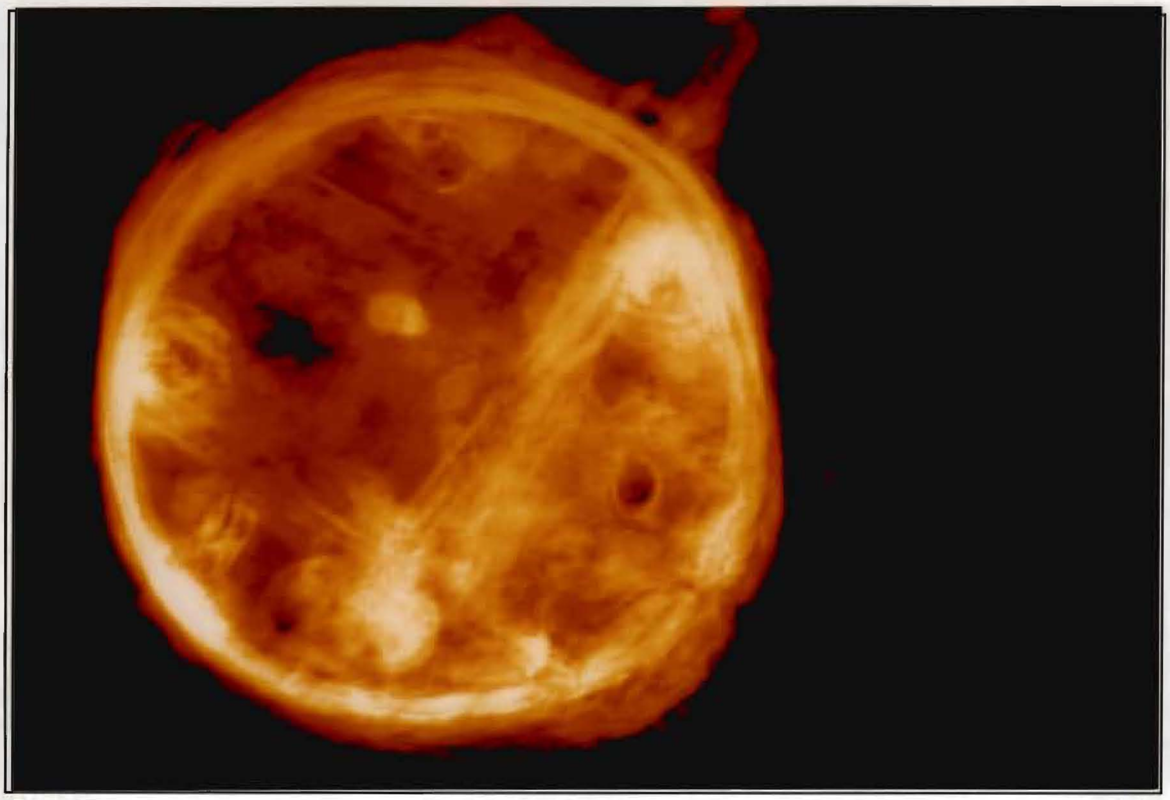

Figure 48 *An example of enhancement using the heated object LUT shown in figure 10.*

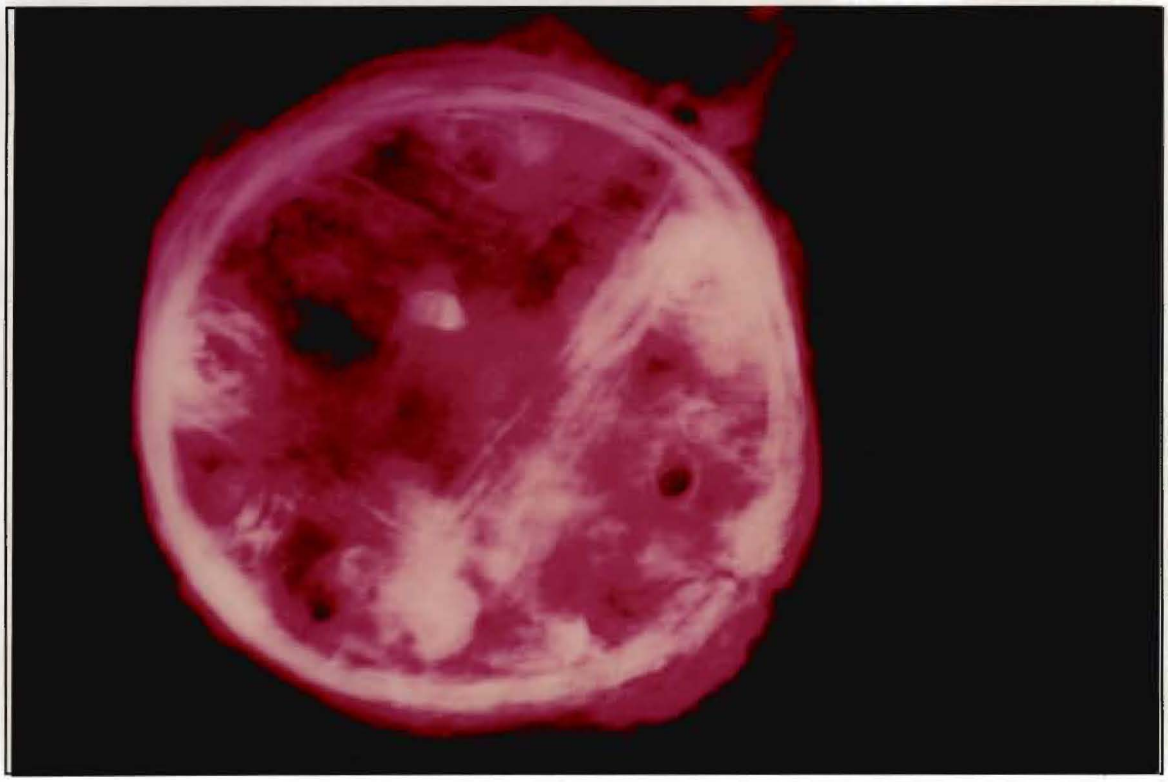

Figure 49 *Enhancement using the magenta LUT shown in figure 14.*

Figure 50 *Application of a gradient filter working in an easterly direction.*

Figure 51 *The effects of a Laplacian filter. Note that the areas of poor digitisation i.e. the vertical lines are particularly visible.*

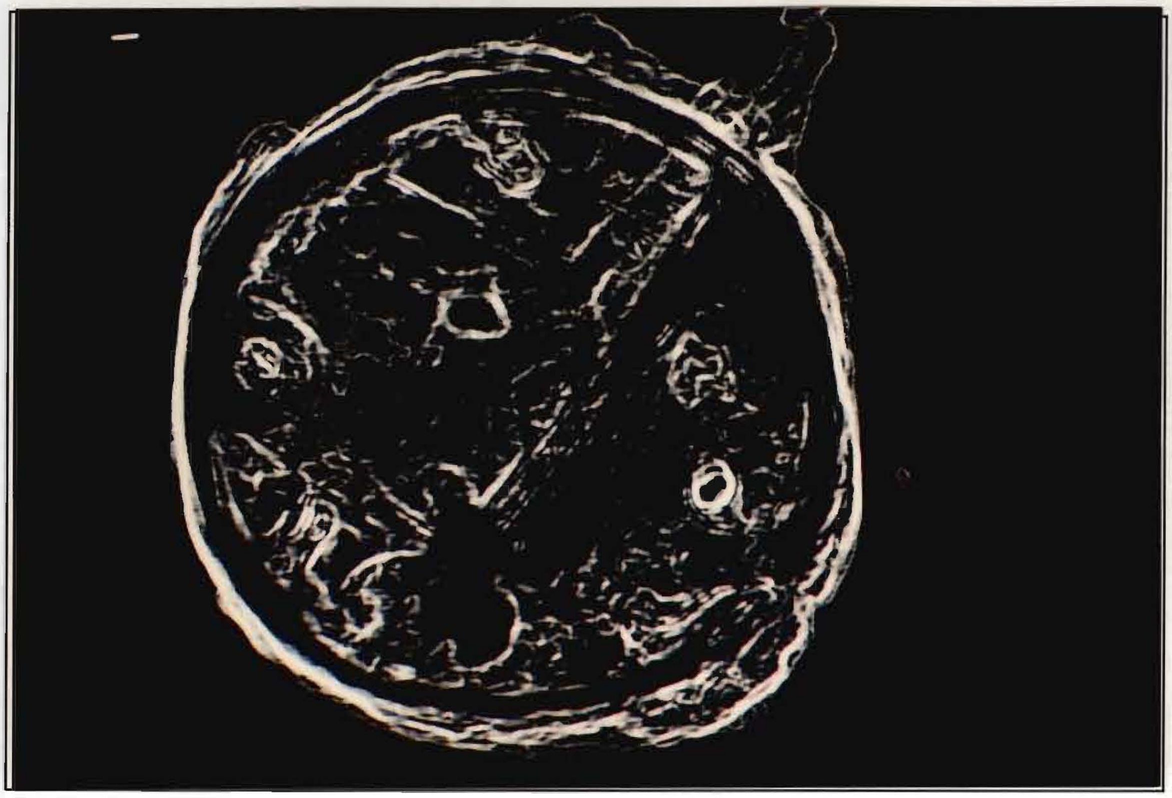

Figure 52 *The effect of a Sobel filter on the original radiograph. Historgram stretching was applied in order to extend the contrast of the image.*

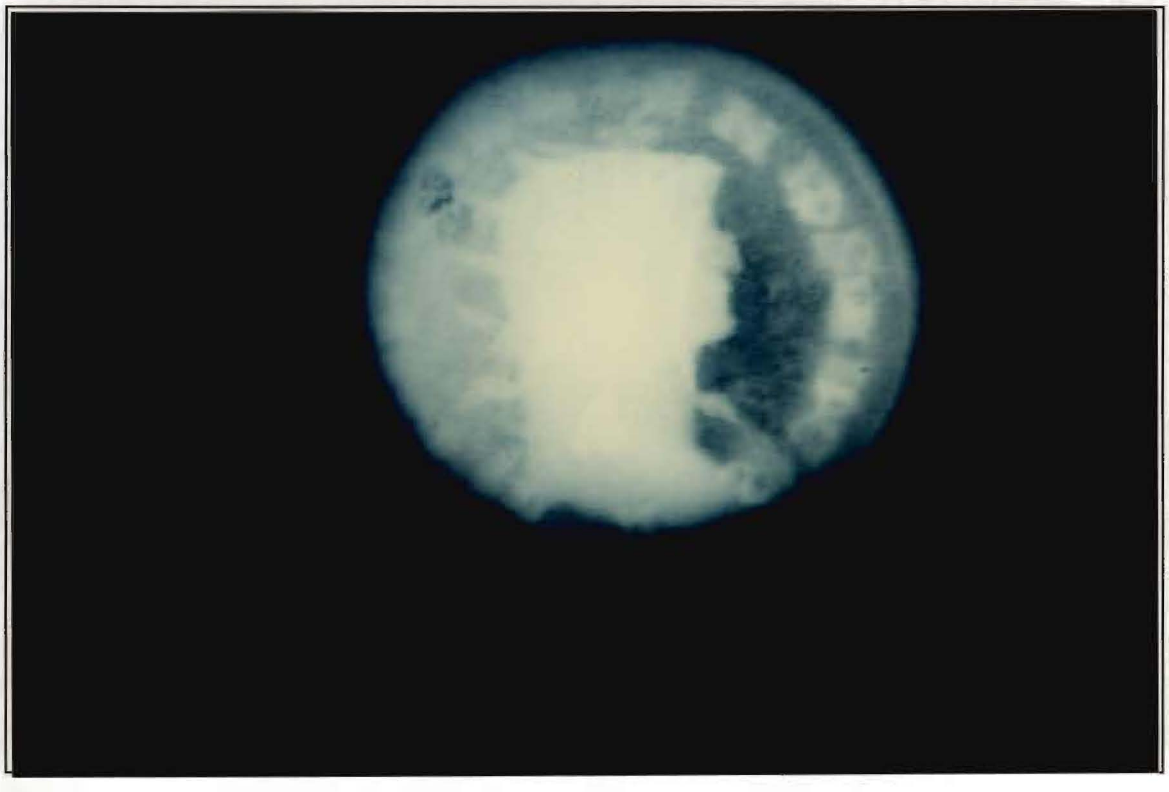

Figure 53 *Original radiograph of the copper alloy coin shown in figure 26. The difference in detail is due to the different recording media.*

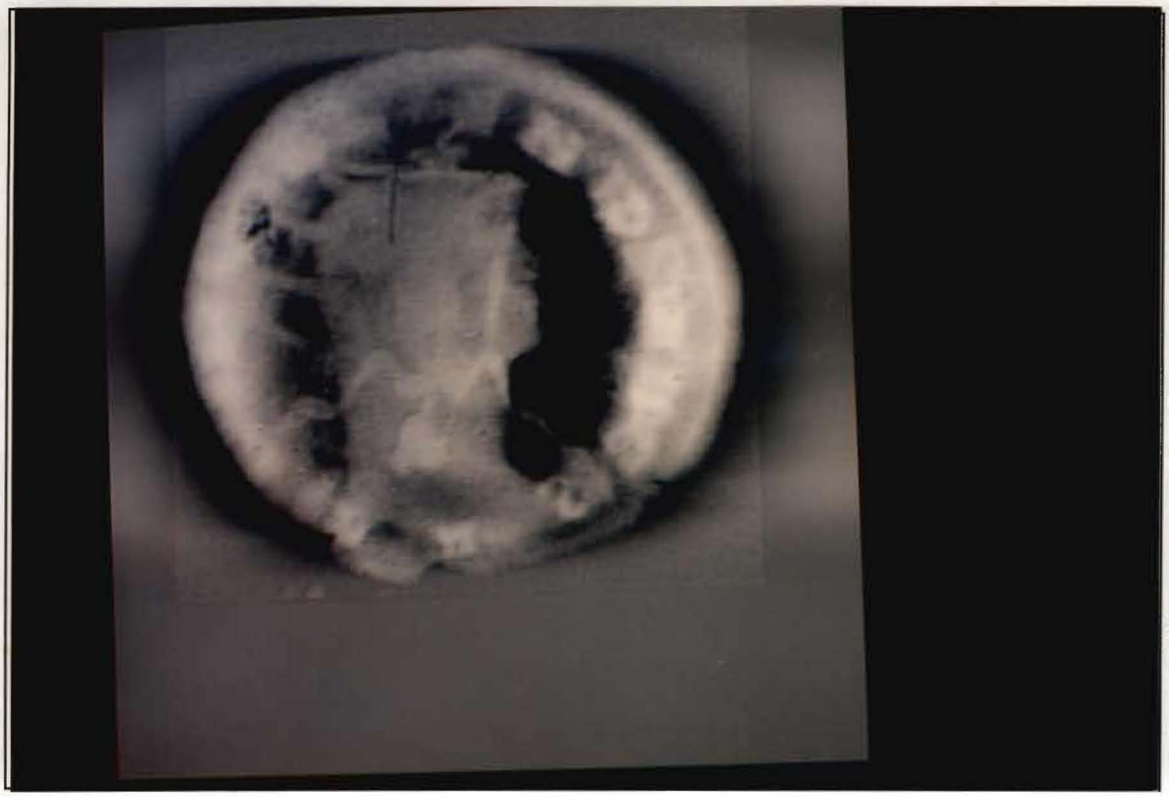

Figure 54 *The coin image after application of FFT to filter out the low frequency component. More detail is visible particularly at the lower edge of the coin.*

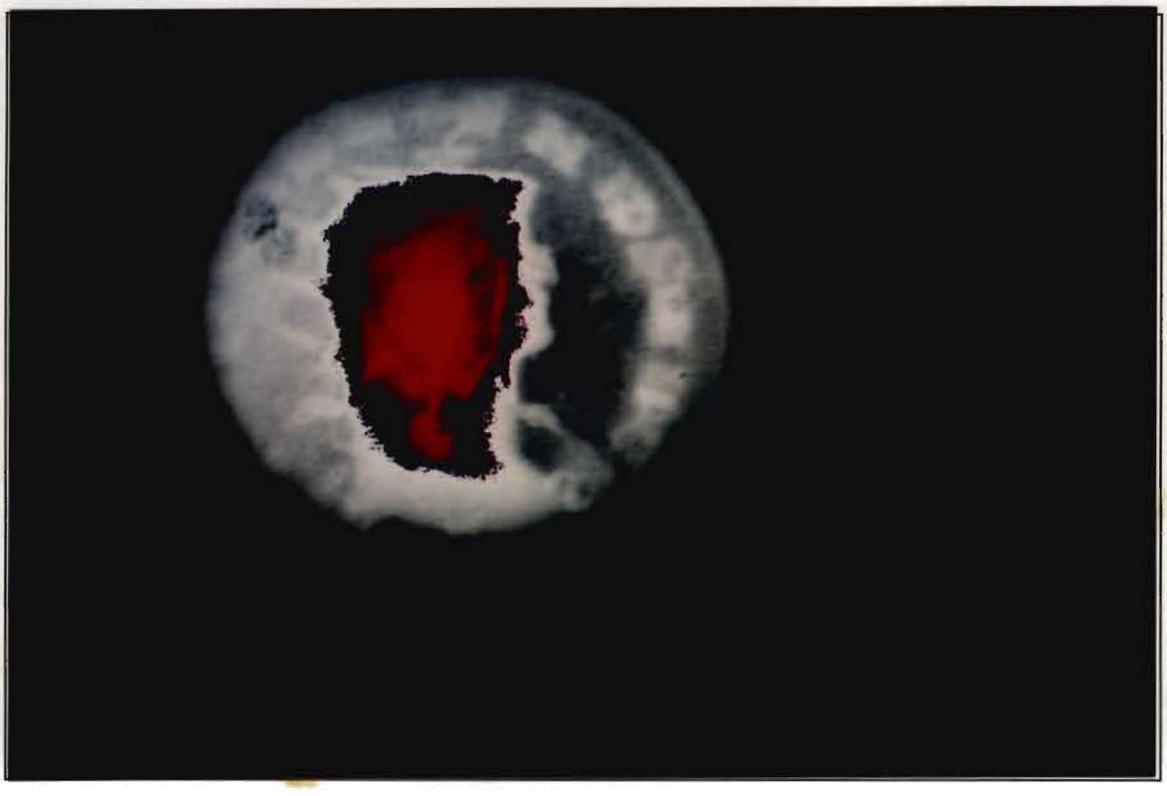

Figure 55 *The central figure on the coin (upside down) highlighted in shades of red using the LUT in figure 11.*

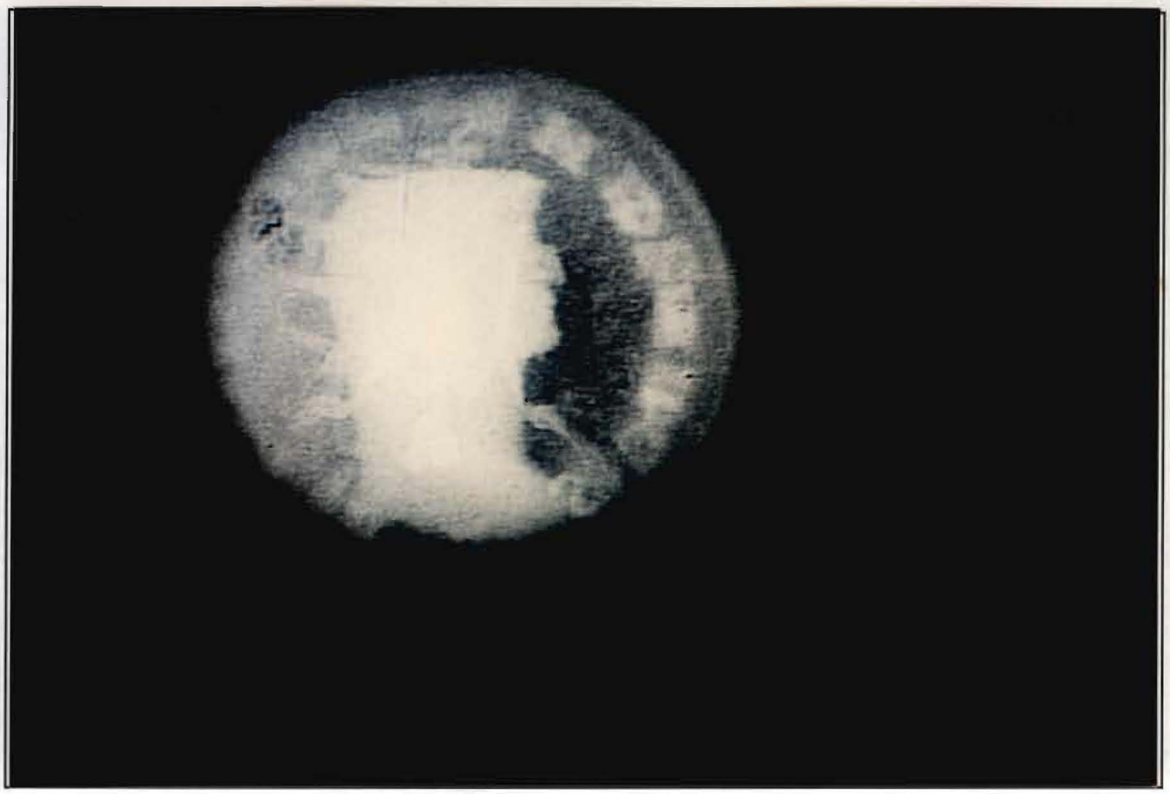

Figure 56 *The radiograph of the coin after application of a sharpening filter . Note the grainy appearance of the image perhaps reflecting the internal structure of the alloy.*

Figure 57 *The textured surface appearance was produced using a specific gradient filter and contrast adjustment.*

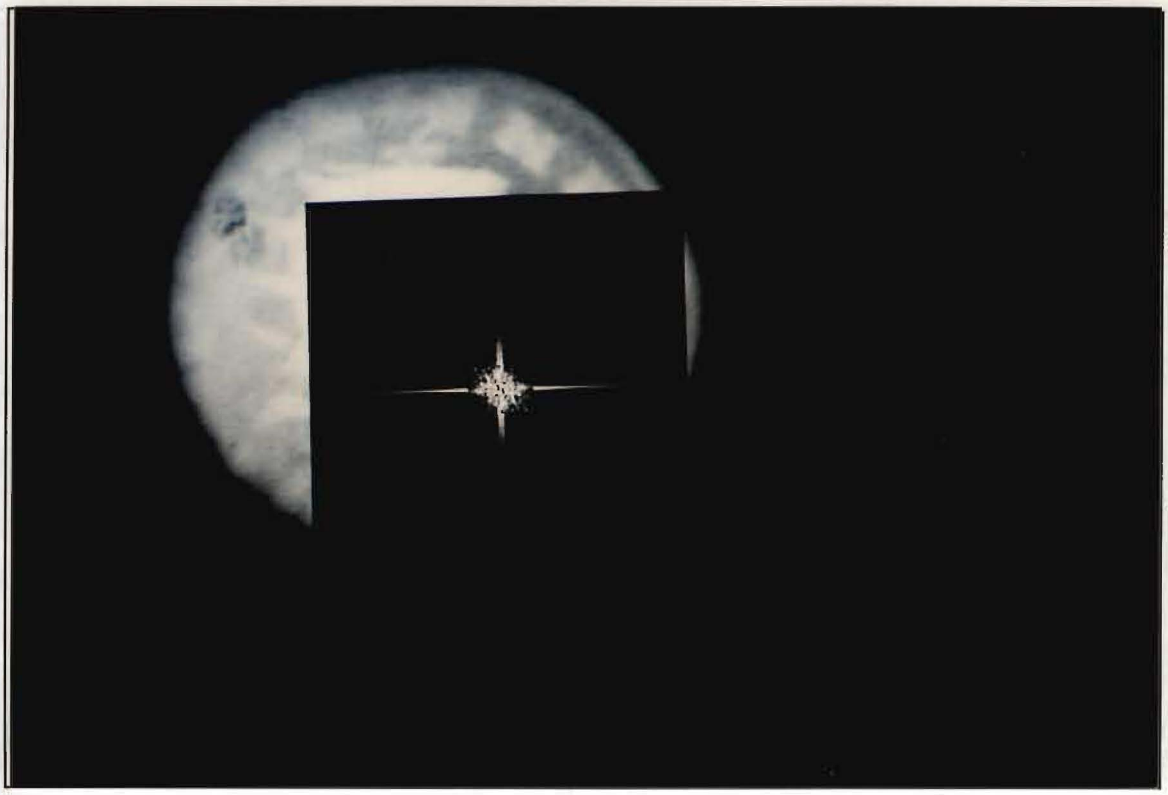

Figure 58 *The energy spectrum of the frequency components within the image of the coin.*

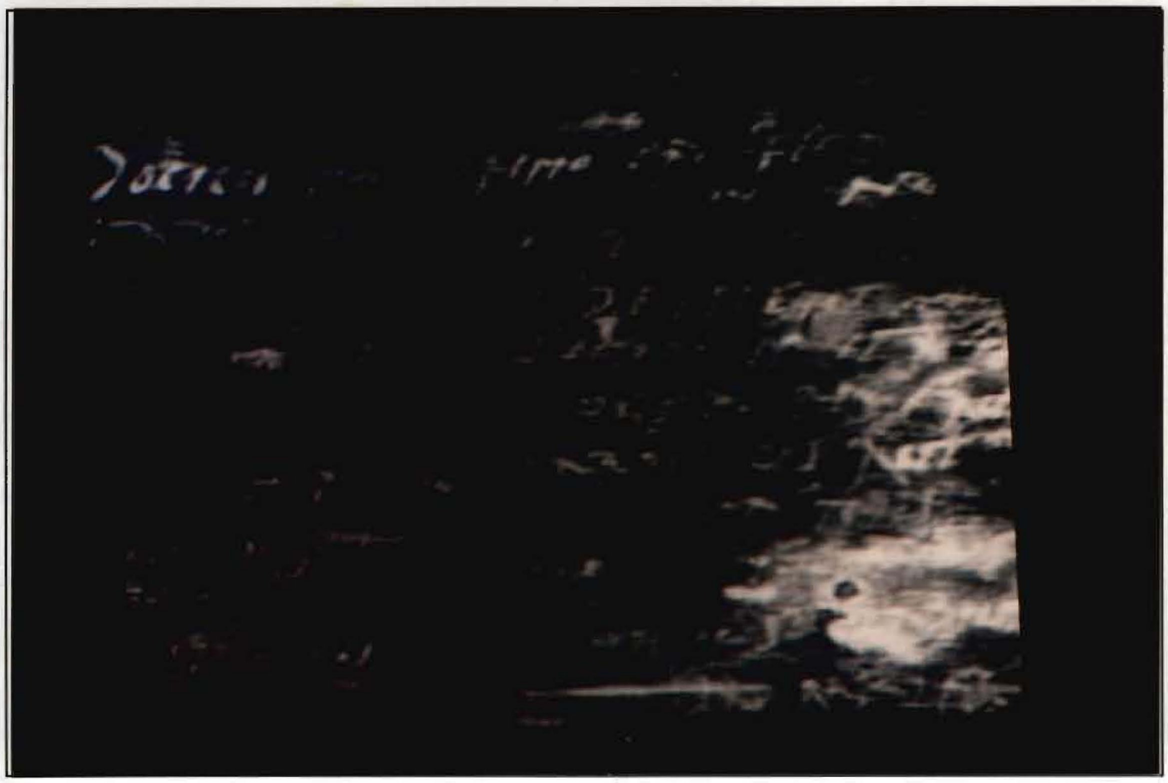

Figure 59 *Infra red photograph of a wooden ink writing tablet. The writing is not visible under normal conditions.*

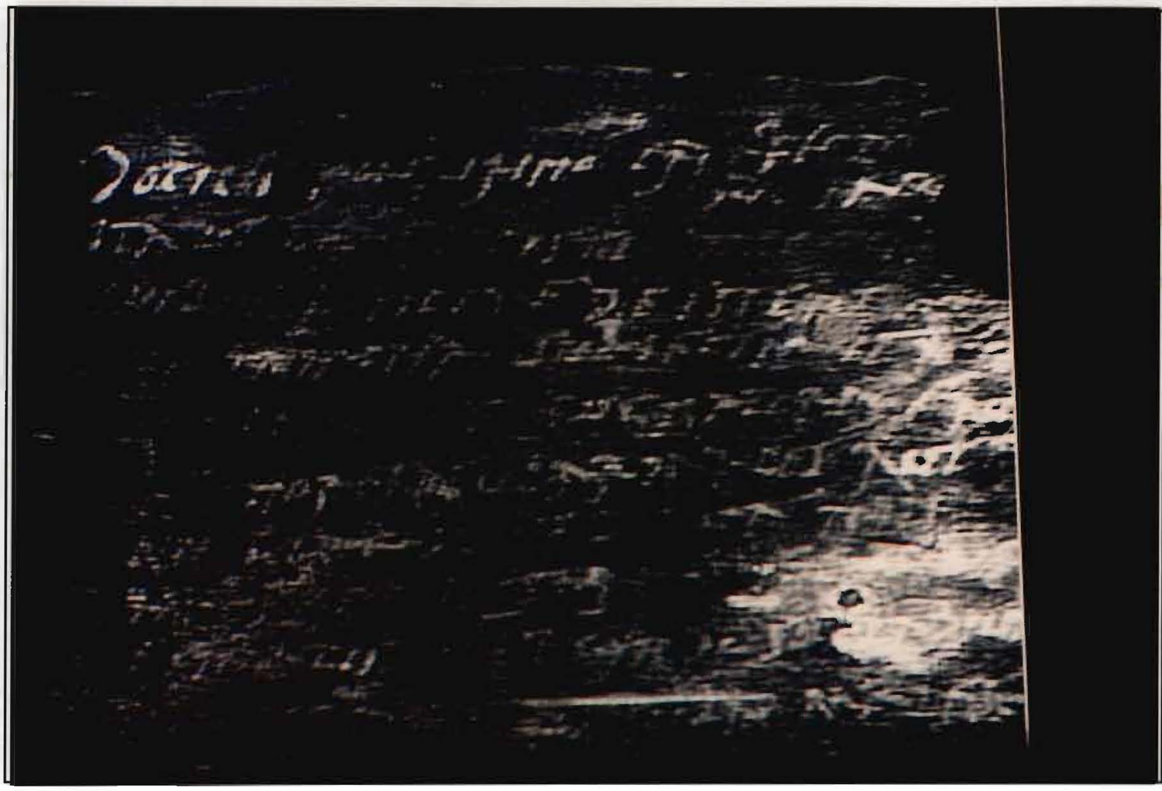

Figure 60 *The writing tablet after application of a simple sharpening filter. Further processing in order to balance the distribution of the grey levels would increase the readability.*

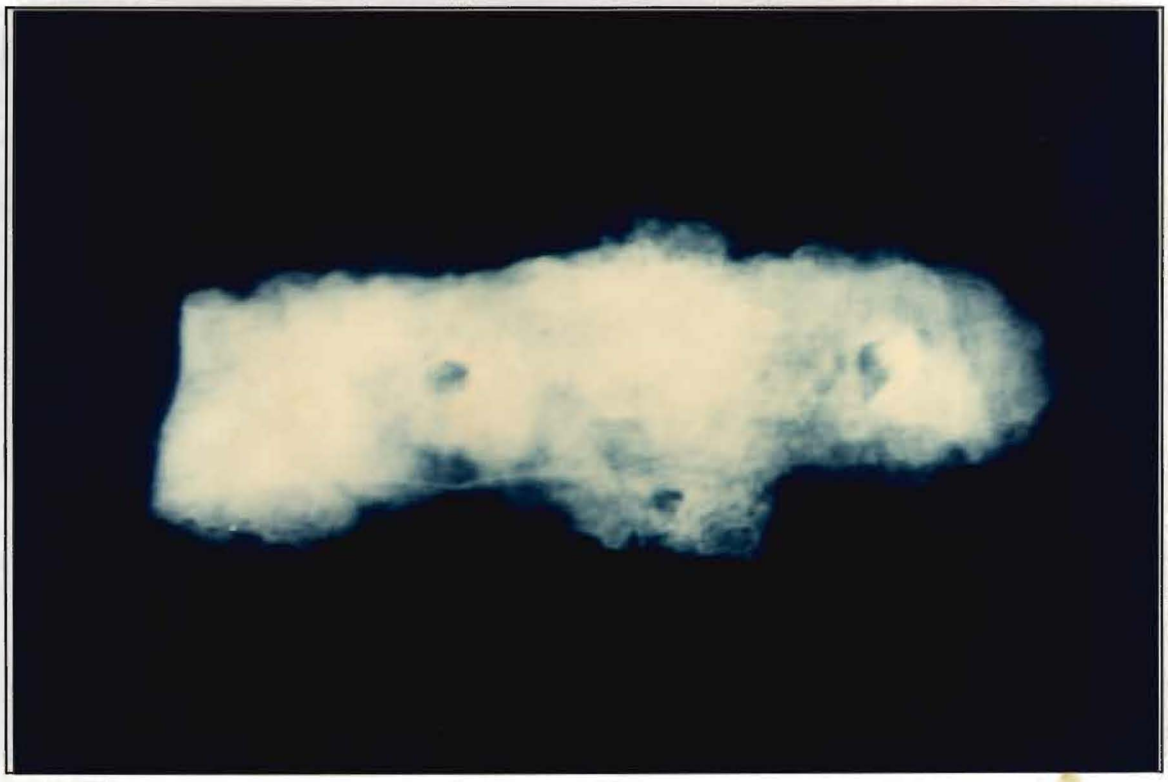

Figure 61 *Original radiograph of an iron coffin fitting. An area of iron replaced wood is just visible on the complete terminal.*

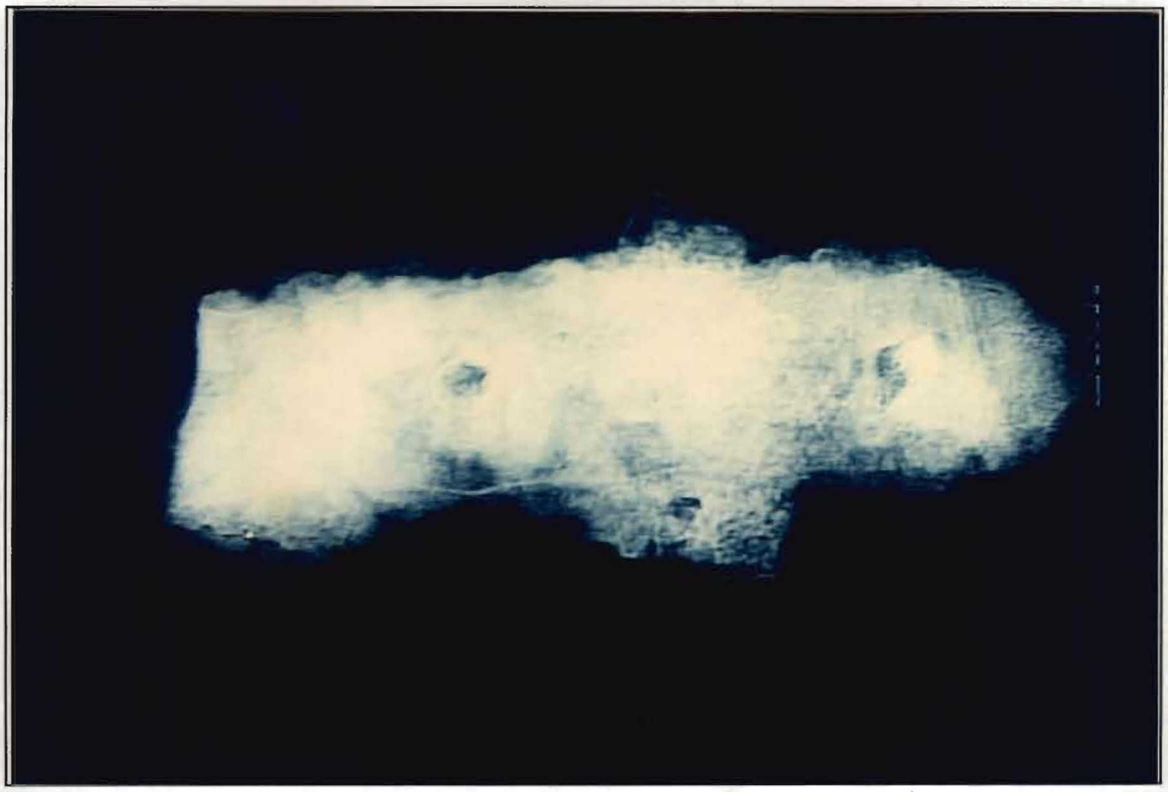

Figure 62 *Application of a high pass filter in order to sharpen the image results in a much clearer impression of the replaced wood..*

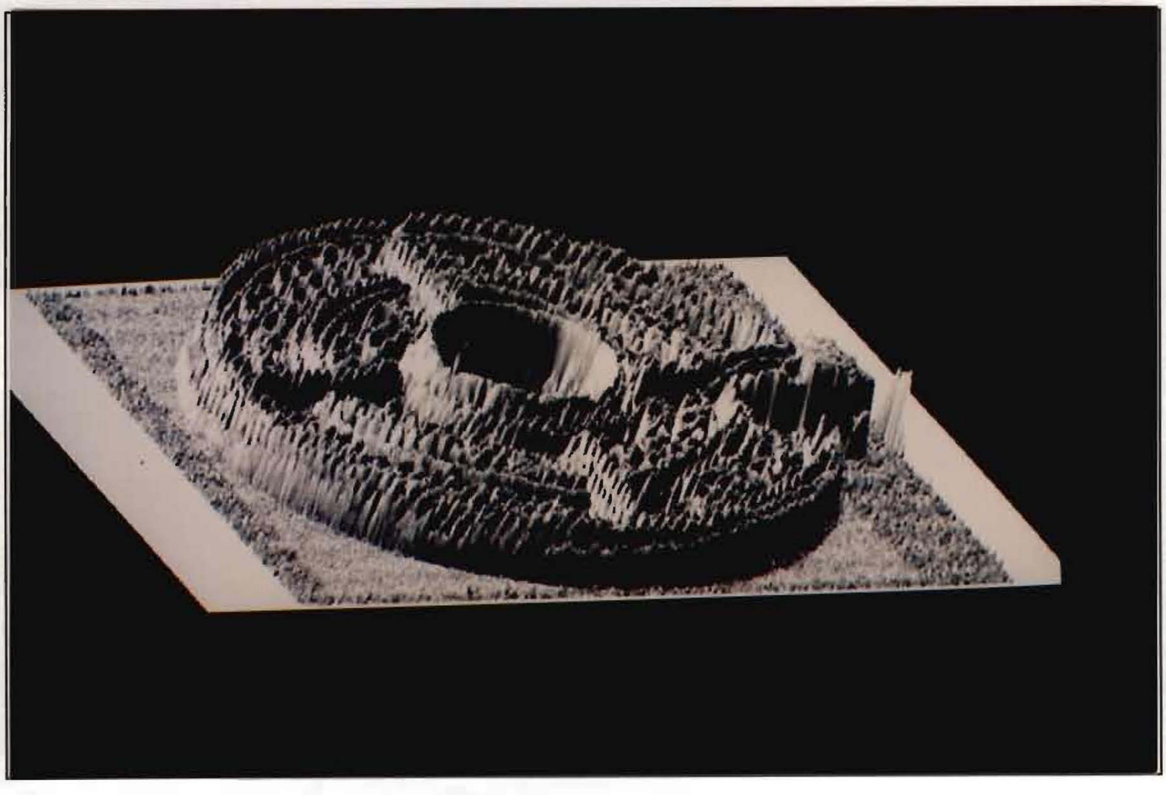

Figure 63 *A 3-D representation of the pendant radiograph in figure 31. The figure illustrates the use of geometrical transformation function.*

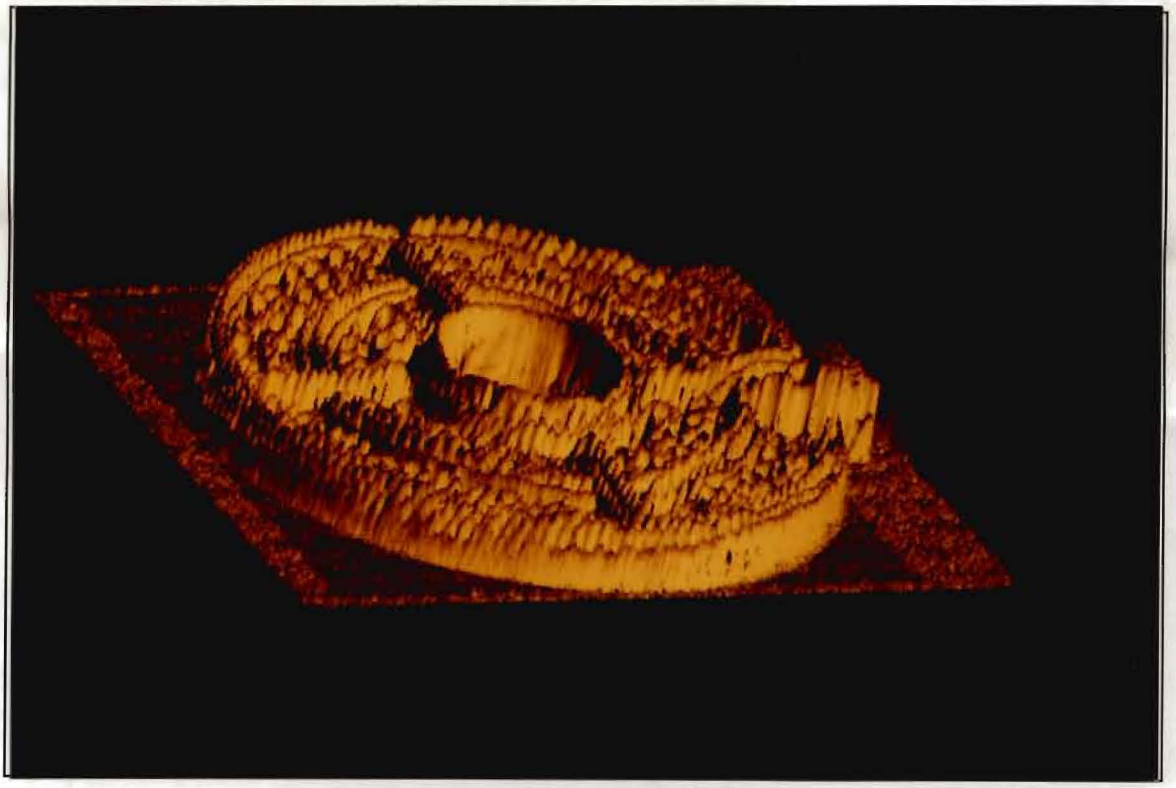

Figure 64 *As in figure 63 with the addition of an pseudo colour LUT based on the heated object scale.*