

High Resolution Colorimetric Imaging of Paintings

Kirk Martinez[‡], John Cupitt[†] and David Saunders[†]

[‡]History of Art Department, Birkbeck College, 43 Gordon Square, London WC1H 0PD

[†]The National Gallery, Trafalgar Square, London WC2N 5DN

ABSTRACT

With the aim of providing a digital electronic replacement for conventional photography of paintings, a scanner has been constructed based on a 3000 x 2300 pel resolution camera which is moved precisely over a 1 metre square area. Successive patches are assembled to form a mosaic which covers the whole area at c. 20 pels/mm resolution, which is sufficient to resolve the surface texture, particularly craquelure. To provide high colour accuracy, a set of seven broad-band interference filters are used to cover the visible spectrum. A calibration procedure based upon a least-mean-squares fit to the colour of patches from a Macbeth Colorchecker chart yields an average colour accuracy of better than 3 units in the CMC uniform colour space.

This work was mainly carried out as part of the VASARI project funded by the European Commission's ESPRIT programme, involving companies and galleries from around Europe. The system is being used to record images for conservation research, for archival purposes and to assist in computer-aided learning in the field of art history. The paper will describe the overall system design, including the selection of the various hardware components and the design of controlling software. The theoretical basis for the colour calibration methodology is described as well as the software for its practical implementation. The mosaic assembly procedure and some of the associated image processing routines developed are described. Preliminary results from the research will be presented.

1. INTRODUCTION

It has been known for some time that a digital replacement was needed for conventional photography. Current documentation techniques in museums and galleries involve the use of transparencies up to 10 x 8" format, which can neither be accurate or stable. The ESPRIT II research programme of the European Community funded the VASARI project, an abbreviation for Visual Arts System for Archiving and Retrieval of Images. This is a particularly suitable acronym since Giorgio Vasari was one of the earliest art historians. The VASARI project aimed to build a new high resolution, colorimetric image capture system capable of competing with photographic resolution while being far more accurate^{1,2}. These images would be used in the same way as transparencies for publishing, research and education but also provide the data for conservation and scientific analysis³. An area of particular interest was colour⁴ and surface texture change⁵. Previously colour measurement was carried out with spot readings or a 256 by 256 pel CID camera⁶. Similarly surface cracks (craquelure) were recorded by photographic details. With either technique it was difficult to make further measurements for later comparison. The VASARI system promised to revolutionise imaging in galleries and museums.

2. APPROACH

It was decided that a resolution would be chosen related to the painting size: ie in pels per millimetre of painting. This meant that one single camera could not be used because the highest resolution available would be around 8 x 8kpels using a linescan CCD. An analysis of paintings showed that a resolution of 20 pels/mm would suffice to detect cracks which were of the order of 0.1mm wide, although microcracks exist which are even smaller. This dictated a resolution of around 20 x 20k for a 1m² painting. Thus some form of positioning device was needed to scan a sensor over the painting area. An area CCD camera was chosen over linescan CCDs to reduce the specifications of the positioning system. This means that a mosaicing technique is required to join together all the sub-frames. It was assumed that the positioning was not exact, so a slight overlap in the sub-frames would be used to calculate the error and join them correctly.




Figure 1 The Scanner System

3. THE SYSTEM

3.1. Overall design

The whole system was to use standard commercial components as far as possible and be relatively inexpensive even though it was the first prototype. The repositioning system is shown Figure 1. The camera is mounted on a vertical portal which has motorised X and Y axes. It was too dangerous to support the paintings in a non-vertical position. Behind the camera is the light projector containing the filters, which illuminates a small patch of painting via fibre-optic cables. All mains power to the light camera and computer is stabilised to avoid problems caused by fluctuating mains. Various calibration charts are mounted on the easel with the painting.

3.2. The Camera

A low resolution camera could not be used if images were to be mosaiced to make a 20 x 20k result, because joining errors would accumulate and too many steps would be required. A high resolution commercially available camera was selected to produce a high resolution image for each exposure. The Kontron ProgRes 3000, was selected for the VASARI acquisition system⁷. This is a commercial development of a camera invented at the Technical University of Munich which offers a resolution of up to 2320 pels. The A/D converter is built into the camera head so the signal can safely be transmitted along the scanner cable to the interface card in the PC. The camera currently in use has an eight-bit A/D converter but a twelve-bit version is available and is under investigation. The sensor within the camera is a standard TV camera CCD which has been masked so as to reduce the size of the active sites. Piezo-crystal actuators mounted around the sensor can be controlled so as to shift the chip by a fraction of the active site separation. The maximum number of displacements in the horizontal and vertical directions is 8 respectively giving the increased resolution from 500 x 290 pels.




Figure 1 The Scanner System

3. THE SYSTEM

3.1. Overall design

The whole system was to use standard commercial components as far as possible and be relatively inexpensive even though it was the first prototype. The repositioning system is shown Figure 1. The camera is mounted on a vertical portal which has motorised X and Y axes. It was too dangerous to support the paintings in a non-vertical position. Behind the camera is the light projector containing the filters, which illuminates a small patch of painting via fibre-optic cables. All mains power to the lights, camera and computer is stabilised to avoid problems caused by fluctuating mains. Various calibration charts are mounted on the easel with the painting.

3.2. The Camera

A low resolution camera could not be used if images were to be mosaiced to make a 20 x 20k result, because joining errors would accumulate and too many steps would be required. A high resolution commercially available camera was selected to give a high resolution image for each exposure. The Kontron ProgRes 3000, was selected for the VASARI acquisition system⁷. This is a commercial development of a camera invented at the Technical University of Munich which offers a resolution of up to 3000 by 2320 pels. The A/D converter is built into the camera head so the signal can safely be transmitted along the scanner ducts to the interface card in the PC. The camera currently in use has an eight-bit A/D converter but a twelve-bit version is available and is under investigation. The sensor within the camera is a standard TV camera CCD which has been masked so as to reduce the size of the active sites. Piezo-crystal actuators mounted around the sensor can be controlled so as to shift the chip by a fraction of the active site separation. The maximum number of displacements in the horizontal and vertical directions is 6 and 8 respectively giving the increased resolution from 500 x 290 pels.

The colour version of the camera has red, green and blue lacquer stripes laid over the sensor. But because the colours of these lacquers are inhomogeneous they are not suitable for accurate colour measurement. The camera in the VASARI system uses the monochrome sensor and an alternative method of colour separation, described below. It was necessary to remove the infrared blocking filter built into the camera, since it absorbed a considerable portion of the red region of the visible spectrum. This required infrared to be filtered out in the light system.

While it is possible to correct geometric distortions in an image, this process is computationally expensive and inevitably results in some loss of image quality. A high-quality lens, which produces little or no distortion, is therefore desirable and a Zeiss Z-Planar f4/32 lens was selected. This lens is designed for producing microfiche films. It has low geometric distortion, low radiometric distortion (vignetting or shading) and a good modulation transfer function, the most important properties demanded by this application. It does however have a maximum aperture of f4, so plenty of light is needed.

Initially the camera control board operated in a PC. Images were grabbed from the camera into a seven megabyte RAM disc and transferred to the workstation over Ethernet. This transfer was the rate-determining step in the acquisition process. An SBUS board which will make it possible to connect the camera directly to the SUN is now used, which gives a frame-grab time of around 6 seconds per channel.

3.3. Positioning System

It is necessary to create a series of sub-images covering the surface of larger paintings. The positioning equipment is used to move the camera and light projector parallel to the plane of the painting. This equipment was built to specification by Time and Precision Ltd (UK). It consists of a rigid steel base mounted on a concrete floor with vibration damping blocks. On the base are two 2.5m stainless steel rails on which the main portal moves. The portal is roughly positioned by hand at an appropriate distance from the painting along these rails and is then locked in place. The portal has two vertical rods between which the horizontal axis is mounted. The camera and light projector are attached to a platform which rests on this horizontal axis. Both axes are motorized and can be moved under computer control by up to 1m in the horizontal and vertical directions to a resolution of 5 μ m. At the far end of the base is the easel upon which paintings are mounted. The camera platform is also motorized, which allows the camera to be moved backwards and forwards about 10cm, which is used for automatic focusing. The stepper-motors are driven by a microprocessor-controlled interface, which is in turn controlled across a serial link by the workstation. The motor interface has two spare output transistors which are used by the workstation to control the filter changer described below.

The accuracy of the positioning equipment was assessed by TÜV-Bayern, a German standards organisation. The repositioning accuracy was found to be better than 10 μ m, with an absolute positioning accuracy of about 30 μ m. When the positioning system was specified, the accuracy necessary for the automatic construction of image mosaics was unknown; with hindsight, perhaps slightly lower positioning accuracy would also produce acceptable results.

To ensure geometric alignment between adjacent sub-images it is necessary to ensure that the focal axis of the camera is perpendicular to the plane of the easel, and that the movement of both the horizontal and vertical axes is parallel to the easel plane. The final alignment of the camera was achieved using a series of very fine spacers, placed beneath the mounting points at the corners of the camera. The tests conducted by TÜV-Bayern indicated the co-planarity of the easel and portal axes. These findings were confirmed by imaging a large, accurate, grid pattern. When the camera was displaced in either the horizontal or vertical direction an error of less than one pixel (that is, one part in 3000) was noted.

3.4. Light System

Unfortunately, the interference filters used are of significant thickness; even a fractional change in the alignment of the filter produces a significant deviation in the light path and leads to alignment errors. Thus they could not be used in front of the camera lens. The solution was to place the interference filters in the optical path between the light source and the fibre-optic guide. This reduced the exposure of the painting as well as overcoming the problem of image registration. The filters were mounted on a wheel which was rotated by a stepper-motor. A small controller unit allowed the filters to be changed under computer control. The light projector is mounted at the rear of the camera platform (Figure 2). The light projector and the fibre-

optic light guides which illuminate the painting move with the camera. This is useful because the light distribution is identical for each sub-image and the same light-distribution correction can be applied to each sub-image.

Figure 3 shows the present light projector schematically. Two 24V, 250W DC tungsten halogen lamps are contained in separate enclosures. Each has a set of collimating optics and an infrared blocking filter. The light from these two sources is fed through a 'Y' shaped fibre-optic guide to the enclosure containing the filters. Thus, the filters are separated from the bulbs, preventing any conductive heating, while the infrared output of the light source is attenuated both by the filter in the light box and by the connecting fibre-optic guide.




Figure 2: The lighting system

In the filter box the light passes through one of seven filters which are secured inside an aluminium wheel. A fan passes a stream of air through channels cut in the filter wheel, cooling the filter in use. Since the sensitivity of the camera varies greatly across the range of frequencies covered by the filters, some of the filters are paired with a wire neutral density filter. The filter wheel is turned by a stepper motor. A microswitch senses the position of notches cut in the edge of the wheel, stopping at each filter. Even with so much light the low transmittance of the filters leads to high noise in the extreme red (700nm) and blue (400nm) images. The workstation (via the stepper-motor interface) can pulse the motor on, to select filters. From the filter box, the light passes into a second fibre-optic guide which divides into six 'tails' each terminated by a frosted lens unit. These provide fairly even illumination over the region to be imaged but it is a time-consuming task to configure the light guides properly.

3.5. Computer System

The camera, positioning equipment and light projector are controlled by a Sun SPARCstation 2GS workstation (currently upgraded to Sparcstation 10/42) which is also used to gather and process the images from the camera. This machine has 32 MBytes of RAM and a 24 bit graphics system. An extra 4GBytes of hard disc storage has been added to provide adequate work space. A Maxtor Tahiti 1GByte optical disc drive is used to archive both the raw camera data and complete calibrated images, but gives a disappointing 150kBytes/s read rate. The software is entirely written in C using the Motif toolkit, except an icon-based human interface (SIAM) developed by Thomson which is based on C++.

optic light guides which illuminate the painting move with the camera. This is useful because the light distribution is identical for each sub-image and the same light-distribution correction can be applied to each sub-image.

Figure 3 shows the present light projector schematically. Two 24V, 250W DC tungsten halogen lamps are contained in separate enclosures. Each has a set of collimating optics and an infrared blocking filter. The light from these two sources is fed through a 'Y' shaped fibre-optic guide to the enclosure containing the filters. Thus, the filters are separated from the bulbs, preventing any conductive heating, while the infrared output of the light source is attenuated both by the filter in the light box and by the connecting fibre-optic guide.




Figure 2: The lighting system

In the filter box the light passes through one of seven filters which are secured inside an aluminium wheel. A fan passes a stream of air through channels cut in the filter wheel, cooling the filter in use. Since the sensitivity of the camera varies greatly across the range of frequencies covered by the filters, some of the filters are paired with a wire neutral density filter. The filter wheel is turned by a stepper motor. A microswitch senses the position of notches cut in the edge of the wheel, stopping at each filter. Even with so much light the low transmittance of the filters leads to high noise in the extreme red (700nm) and blue (400nm) images. The workstation (via the stepper-motor interface) can pulse the motor on, to select filters. From the filter box, the light passes into a second fibre-optic guide which divides into six 'tails' each terminated by a frosted lens unit. These provide fairly even illumination over the region to be imaged but it is a time-consuming task to configure the light guides properly.

3.5. Computer System

The camera, positioning equipment and light projector are controlled by a Sun SPARCstation 2GS workstation (currently upgraded to Sparcstation 10/42) which is also used to gather and process the images from the camera. This machine has 32 MBytes of RAM and a 24 bit graphics system. An extra 4GBytes of hard disc storage has been added to provide adequate work space. A Maxtor Tahiti 1GByte optical disc drive is used to archive both the raw camera data and complete calibrated images, but gives a disappointing 150kBytes/s read rate. The software is entirely written in C using the Motif toolkit, except an icon-based human interface (SIAM) developed by Thomson which is based on C++.

4. COLOUR MEASUREMENT

4.1. The Filter Set

One aim in the development of the present image processing system was to record colour information for a whole painting in a generally recognized colour space. A study of the colour separation technique required suggested that up to twelve narrow-band filters would be needed for spectral reconstruction⁸. However a set of seven broad-band filters were selected covering the visible spectrum. It was assumed that a linear combination of their responses (modified by that of the camera and optics) was related to the XYZ tristimulus values⁹. It was initially hoped that the reflectance spectrum could be reconstructed but this is extremely computationally intensive, involves much more storage and would usually be converted to XYZ tristimulus values anyway. The selected set of seven broad band interference selected have roughly Gaussian characteristics with a bandwidth of 70nm at half the maximum transmittance and their peak transmittances range from 400 to 700nm in steps of 50nm. Their transmittances cover the visible spectrum with considerable overlap. Hence, most of the spectral information will be contained in the seven channel image recorded. Although the colour accuracy might be improved slightly by using a greater number of filters, there are advantages in using seven rather than twelve filters (suggested by work carried out at by our partners in Telecom Paris⁸). The time taken to record twelve images would be proportionally longer and finally the data generated by recording twelve colour separation images would be proportionally more.

4.2. Colour Coordinates

It would be possible to record the reconstructed spectrum for each pixel, but the storage requirements for such a digital image would be unrealistically large. It was therefore decided that the data should be converted to and stored in a recognized colour notation. The colour system chosen is the Commission Internationale de l'Eclairage (CIE) XYZ standard colour space¹⁰. Since other CIE colour coordinates, for instance $L^*a^*b^*$, may be derived from the values of XYZ, the latter seem an obvious choice. The CIE $L^*a^*b^*$ colour coordinates are derived from XYZ as follows:

For a given pixel, the tristimulus values (XYZ) can be derived from the camera response through each filter ($\chi_1, \chi_2, \dots, \chi_7$) by matrix multiplication. Thus:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = F \cdot \begin{pmatrix} \chi_1 \\ \chi_2 \\ \cdot \\ \cdot \\ \cdot \\ \chi_7 \end{pmatrix} \quad \text{where} \quad F = \begin{pmatrix} f_{x1} & f_{x2} & \cdot & \cdot & f_{x7} \\ f_{y1} & f_{y2} & \cdot & \cdot & f_{y7} \\ f_{z1} & f_{z2} & \cdot & \cdot & f_{z7} \end{pmatrix}$$

The colour calibration routine described in a later section is designed, therefore, to calculate the values of f_{x1} to f_{z7} in the conversion matrix (F) above. This is achieved by imaging n (where $n \geq 7$) colour standards (with known XYZ values) through each of the filters. The matrix (M) representing the response to each colour through each filter ($\chi_1, \chi_2, \dots, \chi_n$) is combined with the known XYZ values for the n colours ($X_1, Y_1, Z_1, \dots, X_n, Y_n, Z_n$), matrix (K), to generate the least mean square solution for the conversion matrix (F). Thus where

$$F = [(M^T M)^{-1} M^T K]^T$$

where M^T denotes the transpose of matrix M and

$$M = \begin{pmatrix} x_{11} & x_{21} & \dots & x_{71} \\ x_{12} & x_{22} & \dots & x_{72} \\ \vdots & \vdots & \ddots & \vdots \\ x_{1n} & x_{2n} & \dots & x_{7n} \end{pmatrix} \quad K = \begin{pmatrix} X_1 & Y_1 & Z_1 \\ X_2 & Y_2 & Z_2 \\ \vdots & \vdots & \vdots \\ X_n & Y_n & Z_n \end{pmatrix}$$

The conversion matrix is stored by the workstation and used to calibrate the colour of each sub-image acquired by the camera, as described in the section below entitled scanning procedure.

4.3. Colour standards

The present colour reference used is a Macbeth ColorChecker Chart®. This chart comprises twenty four colour patches, including the additive and subtractive primaries and a grey scale. The stability and homogeneity of the colours is good. In contrast to some photographically produced colour control patches provided by film manufacturers, the Macbeth colour patches have smooth spectral curves, similar to those of artists' pigments. The colour of each patch is checked routinely using a reflectance spectrophotometer and these values are used in the calibration.

4.4. Colour difference measurement

A method of colour difference measurement is required for two distinct purposes. First, to assess the accuracy of the colour measurement system and secondly, to assess changes in colour between images recorded at different times. The CIE standard measure of colour difference ΔE^*_{ab} is based on the 1976 $L^*a^*b^*$ colour coordinates (CIELAB colour space)¹⁰. These may be derived from the values of XYZ stored for each pixel. ΔE^*_{ab} is calculated from the L^* , a^* and b^* coordinates of the two colour samples as follows:

$$\Delta E^*_{ab} = (\Delta L^{*2} + \Delta a^{*2} + \Delta b^{*2})^{1/2}$$

where ΔL^* is the difference in L^* for the two samples and so on. A value of ΔE^*_{ab} greater than unity is intended to indicate that the eye can discriminate between the two colours; a 'just perceptible difference'.⁶ Although the CIELAB system is intended to be a uniform colour space there are, unfortunately, non-uniformities. As a result, a ΔE^*_{ab} of two or three for a pair of colours in one region of colour space may not correspond to a perceptible difference, while it may be possible to discriminate between another pair of colours which give a ΔE^*_{ab} of less than one.

A colour difference formula developed by the Society of Dyers and Colourists Colour Measurement Committee (CMC), now adopted as a British Standard, attempts to compensate for the non-uniformity of the CIELAB space¹¹. Using the CMC colour difference measure (ΔE^*_{CMC}) gives a better indication of the colour discriminating ability of the human eye and it is this measure that has been used to quantify colour difference. To calculate CMC values firstly the hue angle, h_{ab} , and chroma, C^*_{ab} , are calculated from a^* and b^* as follows:

$$h_{ab} = \arctan(b^*/a^*) \quad \text{and} \quad C^*_{ab} = (a^{*2} + b^{*2})^{1/2}$$

The hue difference between the two colours is determined from the CIELAB colour difference ΔE^*_{ab} .

$$\Delta H^*_{ab} = (\Delta E^{*2}_{ab} - \Delta L^{*2} - \Delta C^{*2}_{ab})^{1/2}$$

The CMC colour difference (ΔE^*_{CMC}) is then calculated:

$$\Delta E^*_{CMC} = ((\Delta L^*/S_L)^2 + (\Delta C^*_{ab}/S_C)^2 + (\Delta H^*_{ab}/S_H)^2)^{1/2}$$

$$M = \begin{pmatrix} x_{11} & x_{21} & \dots & x_{71} \\ x_{12} & x_{22} & \dots & x_{72} \\ \vdots & \vdots & \ddots & \vdots \\ x_{1n} & x_{2n} & \dots & x_{7n} \end{pmatrix} \quad K = \begin{pmatrix} X_1 & Y_1 & Z_1 \\ X_2 & Y_2 & Z_2 \\ \vdots & \vdots & \vdots \\ X_n & Y_n & Z_n \end{pmatrix}$$

The conversion matrix is stored by the workstation and used to calibrate the colour of each sub-image acquired by the camera, as described in the section below entitled scanning procedure.

4.3. Colour standards

The present colour reference used is a Macbeth ColorChecker Chart®. This chart comprises twenty four colour patches, including the additive and subtractive primaries and a grey scale. The stability and homogeneity of the colours is good. In contrast to some photographically produced colour control patches provided by film manufacturers, the Macbeth colour patches have smooth spectral curves, similar to those of artists' pigments. The colour of each patch is checked routinely using a reflectance spectrophotometer and these values are used in the calibration.

4.4. Colour difference measurement

A method of colour difference measurement is required for two distinct purposes. First, to assess the accuracy of the colour measurement system and secondly, to assess changes in colour between images recorded at different times. The CIE standard measure of colour difference ΔE^*_{ab} is based on the 1976 $L^*a^*b^*$ colour coordinates (CIELAB colour space)¹⁰. These may be derived from the values of XYZ stored for each pixel. ΔE^*_{ab} is calculated from the L^* , a^* and b^* coordinates of the two colour samples as follows:

$$\Delta E^*_{ab} = (\Delta L^{*2} + \Delta a^{*2} + \Delta b^{*2})^{1/2}$$

where ΔL^* is the difference in L^* for the two samples and so on. A value of ΔE^*_{ab} greater than unity is intended to indicate that the eye can discriminate between the two colours; a 'just perceptible difference'.⁶ Although the CIELAB system is intended to be a uniform colour space there are, unfortunately, non-uniformities. As a result, a ΔE^*_{ab} of two or three for a pair of colours in one region of colour space may not correspond to a perceptible difference, while it may be possible to discriminate between another pair of colours which give a ΔE^*_{ab} of less than one.

A colour difference formula developed by the Society of Dyers and Colourists Colour Measurement Committee (CMC), now adopted as a British Standard, attempts to compensate for the non-uniformity of the CIELAB space¹¹. Using the CMC colour difference measure (ΔE^*_{CMC}) gives a better indication of the colour discriminating ability of the human eye and it is this measure that has been used to quantify colour difference. To calculate CMC values firstly the hue angle, h_{ab} , and chroma, C^*_{ab} , are calculated from a^* and b^* as follows:

$$h_{ab} = \arctan(b^*/a^*) \quad \text{and} \quad C^*_{ab} = (a^{*2} + b^{*2})^{1/2}$$

The hue difference between the two colours is determined from the CIELAB colour difference ΔE^*_{ab} .

$$\Delta H^*_{ab} = (\Delta E^{*2}_{ab} - \Delta L^{*2} - \Delta C^{*2}_{ab})^{1/2}$$

The CMC colour difference (ΔE^*_{CMC}) is then calculated:

$$\Delta E^*_{CMC} = ((\Delta L^*/S_L)^2 + (\Delta C^*_{ab}/S_C)^2 + (\Delta H^*_{ab}/S_H)^2)^{1/2}$$

$$\text{where } S_L = 0.511 \quad \text{if } L^* < 16 \\ S_L = 0.040957L^* / (1 + 0.01765L^*) \quad \text{otherwise}$$

$$S_C = 0.0638C^*_{ab} / (1 + 0.0131C^*_{ab}) + 0.638$$

$$S_H = (fT + 1 - f)S_C$$

$$f = (C^*_{ab})^4 / (C^*_{ab})^4 + 1900)^{1/4}$$

$$T = 0.56 + |0.2\cos(h_{ab} + 168)| \quad \text{for } 164^\circ < h_{ab} < 345^\circ$$

$$T = 0.36 + |0.4\cos(h_{ab} + 35)| \quad \text{otherwise}$$

A method of assessing the accuracy of colour measurement for the acquisition of each painting has been developed which is based on the Macbeth chart. The colour for each patch is measured using the CCD and colour separation filters and XYZ are calculated using the conversion matrix. These values are then transformed to CIE $L^*a^*b^*$ as outlined above. The CIE $L^*a^*b^*$ values for the colour patches have previously been measured spectrophotometrically. The difference between the two set of CIELAB data, ΔE^*_{CMC} , is calculated for each colour. It is the average ΔE^*_{CMC} for these representative colours that is used as the measure of colour accuracy. The typical figure at present is $\Delta E^*_{CMC} = 2.3$.

5. SCANNING PROCEDURE

5.1. Acquisition

The first phase is *acquisition*, when raw data is taken from the camera and stored on hard disc. In the second *calibration* phase, the data is corrected and processed. XYZ colour coordinates are calculated for each pixel and an image suitable for display on the workstation's screen is generated. Mounted with the painting on the easel are a set of calibration targets:

- A white reference; a piece of smooth white card, used to measure non-uniformities in light distribution.
- A resolution target; a target of known size used to determine the scanning resolution accurately.
- A colour reference; the Macbeth chart, which is used for colour calibration.
- A grey level reference; a standard Kodak grey scale, which is used to measure and correct non-linearities in the response of the camera's CCD sensor.

The stages in the acquisition procedure are outlined below.

- The painting is mounted on the easel. Ideally, it should be in the same plane as the calibration charts, to within a few millimetres, and perpendicular to the focal axis of the camera. The painting is supported on a foam block and secured in position with movable clamps.
- The portal is positioned by hand to give approximately the required number of pixels per millimetre. To assist in this stage a number of positions, which correspond to standard scanning resolutions, have been marked on a 2.5m rule attached to the base of the equipment.
- The light guides are adjusted to give the most intense illumination possible while maintaining a uniform distribution of light across the area to be imaged.
- The camera is roughly focused by hand, then fine focused automatically using a section of the painting illuminated with 'green' light from the filter which has maximum transmittance at 550nm. The automatic focusing routine uses the small axis in the camera platform to move the camera perpendicular to the plane of the painting. The 'sharpness' of the image is measured (using a pel difference measure) and a position chosen that maximises sharpness.
- The automatic acquisition program, ACQUIRE, is started. The operator enters the position of each calibration target and the position and size of the painting. If any of these locations have been previously entered and stored, they can be recalled from a data file.

The remainder of the acquisition process is completely automatic and ACQUIRE executes a sequence of operations. The scanning resolution is accurately determined; measurements are made automatically from the resolution target described above. The program uses this information to calculate the precise sub-image displacements required to build the final mosaic. The camera gain for each filter is determined. The ProgRes 3000 camera has a computer-controllable amplifier between the sensor and A—D converter. The level of amplification is selected so that the digital output generated completely fills the available range of values. Seven images of the white target are made, one through each filter. These images are used in the calibration phase to correct non-uniformities in the distribution of light. Seven images of the grey scale are made which are used in the calibration phase to correct non-linearities in the response of the CCD. Seven images of the Macbeth chart are acquired. These images are used in the calibration phase to calculate the colour correction matrix described above. A series of precisely overlapping frames covering the surface of the painting is then acquired through each of the seven filters.




Figure 3. CIE 1931 xy plot showing the measured and real positions of the 24 colours of the Macbeth Colorchecker chart.

6. Correction and calibration

The calibration matrix is recalculated for every painting to compensate for system changes. Once acquisition is complete, the automatic calibration program, CALIBRATE, is started. The operator indicates in which directory the raw data from the painting is located and into which directory the results are to be written. CALIBRATE then executes the following sequence of operations and calibrations:

The remainder of the acquisition process is completely automatic and ACQUIRE executes a sequence of operations. The scanning resolution is accurately determined; measurements are made automatically from the resolution target described above. The program uses this information to calculate the precise sub-image displacements required to build the final mosaic. The camera gain for each filter is determined. The ProgRes 3000 camera has a computer-controllable amplifier between the sensor and A—D converter. The level of amplification is selected so that the digital output generated completely fills the available range of values. Seven images of the white target are made, one through each filter. These images are used in the calibration phase to correct non-uniformities in the distribution of light. Seven images of the grey scale are made which are used in the calibration phase to correct non-linearities in the response of the CCD. Seven images of the Macbeth chart are acquired. These images are used in the calibration phase to calculate the colour correction matrix described above. A series of precisely overlapping frames covering the surface of the painting is then acquired through each of the seven filters.




Figure 3. CIE 1931 xy plot showing the measured and real positions of the 24 colours of the Macbeth Colorchecker chart.

6. Correction and calibration

The calibration matrix is recalculated for every painting to compensate for system changes. Once acquisition is complete, the automatic calibration program, CALIBRATE, is started. The operator indicates in which directory the raw data from the camera is located and into which directory the results are to be written. CALIBRATE then executes the following sequence of corrections and calibrations:

- Each image is corrected to account for non-uniformity in the distribution of light using the image of the white target made through the appropriate filter.
- The images of the grey level chart are analysed. The camera response to each grey is compared to its known lightness value. Seven tables are generated which when applied to an image compensate for any non-linearities in the response of the CCD. A separate table is necessary for each filter, since the precise shape of the camera response function varies with the spectral power distribution of the light reaching the sensor and with camera gain. The remaining images are then corrected for non-linearity errors.
- The seven corrected images of the Macbeth chart are analysed using the procedure described earlier, and a colour conversion matrix generated. Usually, all the colour patches bar the black and white are used in determining the solution.
- The Macbeth chart is converted to XYZ colour space and the difference between the actual and measured XYZ calculated. As already described, the average of these difference measurements is used as a guide to the calibration accuracy (see Figure 3).
- Finally, each sub-image is converted to XYZ colour space using the colour conversion matrix.

7. Mosaicing

The separate XYZ sub-images must then be joined in a mosaic to make a single XYZ image of the whole painting. Two adjacent sub-images are displayed on the workstation's screen. The operator selects a point in one of the sub-images, and the corresponding point in the other sub-image. From this information the program estimates the region of overlap between the two sub-images. This region is then divided into three parts. A further 20 points with good contrast are selected automatically in each of the three regions in the first image. Using the operator's initial estimation as a guide, the program searches the second image for points corresponding to the 60 points selected in the first image. The current version of the program can tolerate an error of ± 10 pixels in the initial estimation. This refinement process removes the need for the operator to assess the initial points accurately.

A straight line fit for the 60 points is calculated. After discarding those points which deviate greatly from the straight line, the average of the remaining displacements is computed and used as the offset with which the two images should be joined. In the region of overlap the two images are merged to give a smooth transition. The geometric stability of the positioning equipment and lens is such that it is not necessary to perform any rotation or rescaling of the sub-images during the mosaic assembly procedure. If required, the software can compensate for such distortion by resampling one of the sub-images. The merged image is stored and the process repeated to build strips of the mosaic. Once these strips have been produced, they are joined in the same manner to yield the final image. Automatic mosaicing is also possible using the calculated resolution to give a starting point instead of manually selecting a point.

8. Display

The XYZ images generated by the calibration procedure are converted to a set of RGB values to drive a particular monitor based on its characteristics: a Sony monitor and a Barco Calibrator have been used. To compute this transform it is necessary to know the chromaticities of the three phosphors used in the display tube, the minimum and maximum light output and the gamma of each of the electron guns. The transform from XYZ to rgb calculated from this information will not be detailed here¹².

The gamut of colours that can be generated by a CRT monitor is much smaller than the gamut of colours found in paintings⁹. To obtain an image in which the colours displayed appear as near to the original colours as possible, it is necessary to make a uniform reduction in the size of the gamut of the measured colours so that this 'fits' within the range.

The large size of the final mosaic makes it impossible to view the whole image, even on a high resolution computer monitor. A lower resolution version of the image is created which will fit entirely on the screen. An area of this low resolution image

can be selected with the mouse and that portion retrieved for display from the high resolution image (see Figure 4). The use of virtual memory file mapping makes panning fairly responsive, especially as pages are cached. This has avoided the use of expensive specialised image memory display systems.




Figure 4 Example of image display (The Garvagh Madonna by Raphael)

9. Applications

9.1. General

The VASARI system is being used for a number of research programmes. In order to facilitate this work, a new image processing package (VIPS/IP) has been written which allows non-specialist users to manipulate images on the workstation. Accurate, permanent records are being made for conservation research as an aid to the monitoring of both changes in colour and changes in surface texture. Additionally, by connecting an infrared camera to the system, it has been possible to make quality infrared reflectogram images of the underdrawing in paintings. Such high quality images can be the source for high resolution images or details for public access systems and computer aided teaching and learning systems.

9.2. Colour change measurement

In order to assess possible changes in colour it is necessary to record images of the painting at different times. For an accurate comparison of the measured colour it is essential that the two images can be exactly superimposed. The second time a painting is recorded, the image is compared with the first image to assess any changes in colour.

can be selected with the mouse and that portion retrieved for display from the high resolution image (see Figure 4). The use of virtual memory file mapping makes panning fairly responsive, especially as pages are cached. This has avoided the use of expensive specialised image memory display systems.




Figure 4 Example of image display (The Garvagh Madonna by Raphael)

9. Applications

9.1. General

The VASARI system is being used for a number of research programmes. In order to facilitate this work, a new image processing package (VIPS/IP) has been written which allows non-specialist users to manipulate images on the workstation¹³. Accurate, permanent records are being made for conservation research as an aid to the monitoring of both changes in colour and changes in surface texture. Additionally, by connecting an infrared camera to the system, it has been possible to make high quality infrared reflectogram images of the underdrawing in paintings. Such high quality images can be the source for lower resolution images or details for public access systems and computer aided teaching and learning systems.

9.2. Colour change measurement

In order to assess possible changes in colour it is necessary to record images of the painting at different times. For an accurate comparison of the measured colour it is essential that the two images can be exactly superimposed. The second time a painting

is placed on the easel, the scanning resolution, and perhaps the exact orientation, will not be the same as the first time. The difference in resolution and orientation can be detected using the craquelure analysis routine described in the next section. Unless the recording conditions are identical it will be necessary to resample one of the images to obtain exact registration between them. The first order transformation algorithm to accomplish this step has been developed as part of the mosaic assembly.

While changes in the positioning of the painting may be overcome easily with software, it is far more difficult to compensate for dimensional changes in the painting caused by, for instance, the warping of a panel due to alterations in relative humidity. While controlling the climatic conditions in the laboratory help prevent this, it may not fully compensate for intervening distortions. Although these dimensional changes may render a complete assessment of changes in colour impossible, a comparison of the two images may yield interesting information on the nature of the distortion.

Once a pair of colorimetric images have been acquired, the colour changes can be assessed by calculating a third image in which each pixel represents the ΔE^*_{CMC} between corresponding pixels in the source images. In comparing the images, a filter is used to remove those differences which might arise from signal noise. Only changes above a certain threshold are retained. The high resolution of the image means that it is unlikely that any colour change would be limited to an area of just a few pixels. Isolated points of difference are therefore discarded. The final difference image can be converted to a false colour map for display. Images of this type provide a very simple indication of the areas of the painting which are changing most rapidly. Alternatively, the difference image can be combined with that of the painting to provide a clear indication of the location of change.

Other kinds of colour difference measurement may also be useful in particular situations. Rather than detecting overall colour change, it may be suspected that certain pigments in the painting are changing hue in a particular manner. In this case it will be more informative to calculate colour change in terms of other parameters, for instance hue angle or change in one of the CIE L^* , a^* or b^* parameters. If, for example, it is believed that certain pigments are becoming yellowed with age, the extent to which pixels in the new image are 'yellowier' than the corresponding pixels in the old image might be shown by producing a map showing the increase in CIE b^* value.

9.3. Surface texture analysis

In collaboration with colleagues at the Doerner Institut in Munich, images are being made before and after the transport of paintings to monitor and assess the development of craquelure as paintings are moved on loan. Software has been developed in Munich that extracts the craquelure pattern from images, which are often recorded in raking light. A second image is recorded after the painting has travelled to an exhibition. The craquelure pattern is extracted by digital filtering and morphological analysis. Common features in the extracted pattern of craquelure in the before and after images are used to give superposition. The craquelure patterns are then compared automatically and any differences highlighted in a false-colour map of the painting. A preliminary study in Munich has indicated that the software can detect small changes in surface appearance caused as the result of transporting a painting by road¹⁴.

9.4. Infrared reflectography

The image acquisition and processing software has also been used in the examination of paintings by infrared reflectography. An infrared vidicon camera has been attached to a second workstation via a digitizing board. The infrared sub-images can be corrected and assembled into a mosaic more rapidly and with greater precision than is possible using the conventional photographic method. It is being used in the National Gallery's systematic examination of its early German and Netherlandish collections where it is proving invaluable for the assembly of mosaics comprising up to 360 sub-images¹⁵.

10. CONCLUSIONS AND FUTURE PLANS

The VASARI project showed that accurate high resolution imaging of paintings was possible. As was hoped, the system and its images are being used for a wide range of purposes. Similar scanners are planned at other sites and it is hoped this further development will lead to improvements in price and performance. At present colour images stored on the system cannot easily be reproduced on paper with good colour fidelity. A new project entitled MARC (Methodology for Art Reproduction in Colour)

began in September 1992, and is investigating accurate printing of colorimetric images. These results will be of use to galleries considering establishing in-house desktop publishing facilities for books and catalogues.

11. ACKNOWLEDGEMENTS

This paper presents some of the work carried out as part of European Community's ESPRIT programme. Other institutions involved in VASARI were: The Doerner Institut, Munich (D); Brameur Ltd. (UK); Telecom Paris (F); Thomson-CSF LER, Rennes (F); TÜV-Bayern, Munich (D). The ESPRIT III project MARC involves: Thomson-CSF (LER), The National Gallery, Birkbeck College, The Doerner Institut, Crosfield Ltd (UK), Hirmer Verlag (D), Schwitter (Switzerland).

12. REFERENCES

1. K. Martinez and A. Hamber, 'Towards a colorimetric digital image archive for the visual arts', *Proc. SPIE*, Vol. 1073, pp. 114-121, 1989.
2. K. Martinez, 'High resolution imaging of paintings: the VASARI project', *Microcomputers for information management*, Vol. 8, No. 4, Ablex Publishing Corp., Norwood, New Jersey, December 1991.
3. D. Saunders and J. Cupitt, 'Image Processing at The National Gallery: The VASARI Project', *The National Gallery Technical Bulletin*, Vol. 14, pp. 72-85, 1993.
4. D. R. Saunders, 'The Detection and Measurement of Colour Change in Paintings by Digital Image Processing', *Proc. SPIE*, Vol. 1075, pp. 405-415, 1989.
5. A. Burmester, J. Cupitt, H. Derrien, N. Dessipris, A. Hamber, K. Martinez, M. Müller, D. Saunders, 'The examination of paintings by digital image analysis', *Proc. 3rd International Conference on non-destructive testing, microanalytical methods and environment evaluation for study and conservation of works of art*, Viterbo, pp 201-214, Oct 1992.
6. D. Saunders, 'Colour Change Measurement by Digital Image Processing', *The National Gallery Technical Bulletin*, Vol. 12, pp 66-77, 1988.
7. R. Lenz and U. Lenz, 'ProgRes3000: A Digital Color Camera with a 2D Array CCD Sensor and Programmable Resolution up to 2994 x 2320 Picture Elements', *Proc. SPIE*, Vol. 1357.
8. G. Deconinck, 'Multi-channel Images: Acquisition and Coding Problems' MSc Thesis, Katholieke Universiteit Leuven and Ecole Nationale Supérieure de Télécommunications, 1990.
9. D. Saunders and A. Hamber, 'From pigments to pixels: Measurement and Display of the Colour Gamut of Paintings', *Proc. SPIE*, Vol. 1250, pp 90-102, 1990.
10. Commission Internationale de l'Eclairage, 'Recommendations on uniform color spaces, color difference equations, psychometric color terms', Supplement No.2 to *CIE Publication No.15 (E-2.3.1)*, 1971/(TC-1.3), 1978.
11. F. J. J. Clarke, R. McDonald and B. Rigg, 'Modification to the JCP79 Colour-Difference Formula', *Journal of the Society of Dyers and Colourists*, Vol. 100, pp. 128-132, 1984.
12. D. L. Post, and C. S. Calhoun, 'An Evaluation of Methods for Producing Desired Colors on CRT Monitors', *Color Research and Applications*, Vol. 14, pp. 172-186, 1989.
13. J. Cupitt and K. Martinez, 'An Image Processing System for the Visual Arts', in *Interacting with Images*, Eds. L. MacDonald and J. Vince, Wiley (in Press 1993)
14. A. Burmester and M. Müller, 'The Registration of Transportation Damages using Digital Image Processing', *Zeitschrift für Kunsttechnologie und Konservierung*, Vol. 6, pp. 335-345, 1992.
15. R. Billinge, J. Cupitt, N. Dessipris and D. Saunders, 'A Note on an Improved Procedure for the Rapid Assembly of Infrared Reflectogram Mosaics', *Studies in Conservation*, Vol. 38, in press, 1993.