

National Gallery Technical Bulletin

Volume 14, 1993

NATIONAL GALLERY PUBLICATIONS
LONDON

Series Editor: Ashok Roy

© National Gallery Publications Limited 1993

All rights reserved. No part of this publication may be transmitted in any form or by any means, electronic, or mechanical, including photocopy, recording, or any information storage and retrieval system, without the prior permission in writing from the publisher.

First published in Great Britain in 1993 by
National Gallery Publications Limited
5/6 Pall Mall East, London SW1Y 5BA.

British Library Cataloguing-in-Publication Data
A catalogue record for this journal is available from the British Library.

ISBN 1 85709 033 0
ISSN 0140 7430

Edited by Diana Davies and Emma Shackleton
Design and computer graphics by Neil Aberdeen

Printed and bound in Great Britain by Lawrence-Allen (Colour Printers) Ltd,
Weston-Super-Mare

Front cover: Raphael, *The Garvagh Madonna*, detail.

Picture Credits

Baltimore, Walters Art Gallery: Plate 2, page 22.

Heino, Hannema-De Stuers Foundation: Fig. 1, page 24.

London, British Library: Fig. 3, page 25. By permission of the British Library.

Northumberland, Alnwick Castle, Duke of Northumberland Collection:

Fig. 3, page 10, and Fig. 4, page 11.

Oxford, Ashmolean Museum: Fig. 11, page 15.

Paris, Bibliothèque Nationale: Fig. 2, page 25.

Washington, National Gallery of Art: Fig. 1, page 8, Andrew W. Mellon Collection,

© 1993 National Gallery of Art, Washington.

Windsor Castle, Royal Library, © 1993. Her Majesty The Queen: Fig. 12, page 15.

The publication of this volume of the *National Gallery Technical Bulletin*
has been made possible by the generous support of
Mr and Mrs Frank Richardson of New York

Image Processing at the National Gallery: The VASARI Project

David Saunders and John Cupitt

An earlier volume of this *Bulletin* described the image processing system which had been constructed during the late 1980s to detect and measure changes in the surface colour of paintings.¹ Since that time the Gallery has been fortunate to secure funding from the European Community's ESPRIT II research programme, which has allowed us to build a new image capture and processing system to overcome the deficiencies of the previous equipment. The system described earlier had several shortcomings. First, the General Electric TN2509 camera was of rather low resolution: only 256 by 256 pixels (picture elements). The area covered by one pixel varied with painting size, but for a painting of 1m by 1m would be approximately 4mm by 4mm. The colour within such an area can be far from uniform. Secondly, it was extremely difficult to reposition the camera in order to compare images recorded at different times. Finally, we were concerned that although the colour measurement method was consistent, it was not in line with internationally recognised standards and it would be difficult to compare our results with those obtained by other colour measurement instruments. The new system that has been developed is described in this paper.




Fig. 1 Overall view of the positioning equipment and lighting system.

Equipment

The system in use at the National Gallery at the end of the VASARI project (Visual Arts System for Archiving and Retrieval of Images; see Appendix 1) is shown in Figs. 1 and 2. It comprises a high-resolution monochrome digital camera mounted on a positioning system. Moving with the camera is a light projector containing a set of filters: portions of the painting and the calibration charts are illuminated in turn with each of the filters and a monochrome image taken by the camera. Using the images of the calibration charts as a guide, software in the workstation (main computer) then recombines all these separate monochrome images into a single high-resolution colorimetric image of the painting.

The individual components which make up the system are described in more detail below.

Camera

It was clear that the new system and its image processing capabilities could be used to address a number of conservation-related issues. For image acquisition it was therefore decided that the image should be of sufficiently high resolution to allow the surface texture of the painting to be analysed and should have sufficiently high colour accuracy to allow the detection of colour change. These images would form an archive, against which to compare future measurements in order to yield information on long-term changes in colour or surface appearance. A preliminary analysis of the surface of a variety of paintings indicated that a resolution of between 10 and 20 pixels per millimetre would be a suitable resolution for most applications, but that it might be necessary to increase this to up to about 40 pixels per millimetre for surfaces with extremely fine craquelure.

An image of a painting of size 1m by 1m recorded at a resolution of 10 pixels per millimetre comprises 10,000 by 10,000 pixels. No commercial digital camera can produce such a high-resolution image in one exposure. It is necessary, therefore, to make a series of overlapping sub-images so that the whole surface of the painting can be covered. Software in the workstation then joins these sub-images together to form a single image of the entire painting.

A high-resolution camera, the Kontron ProgRes 3000, was selected for the VASARI acquisition system. The ProgRes 3000 is a commercial development of a camera invented at the Technical University of Munich which offers both very high resolution, some 3000 by 2320 pixels, and excellent geometric stability. In contrast to other cameras of this resolution, the A-D (analogue to digital) converter is built into the camera head rather than contained in a separate unit. This is an advantage since, unlike digital signals, analogue signals deteriorate with transmission distance.

The sensor within the camera is a standard TV camera CCD (charge coupled device) 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. By taking 48 frames, each with the sensor in a slightly different position, the resolution is increased from the standard (500 by 290) to the maximum (3000 by 2320). The so-called Piezo Aperture Displacement (PAD) mechanism is represented schematically in Fig. 3.

Both monochrome and colour versions of the camera are available. The colour camera has red, green and blue lacquer stripes laid over the sensor. Because the colours of these lacquers are inhomogeneous and are thus 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 infra-red blocking filter built into the camera, since it absorbed a considerable portion of the red region of the visible spectrum. A suitable, colourless, infra-red filter was therefore substituted.

The ProgRes 3000 camera was supplied with a seven-bit A-D converter, that is, it could distinguish only 2^7 (128) levels of grey. Although this range was not completely satisfactory, quantisation noise was reduced in the short term by activating the camera's grey level companding circuitry. The companding mode allowed the camera to provide eight-bit accuracy for the first 64 grey levels and six bits for the remaining 192 grey levels. This increased the accuracy of measurement in dark areas of the painting, where quantisation noise is usually most apparent. Recently, the A-D converter has been upgraded to provide eight bits per pixel (256 grey levels). At the time of writing it is hoped that a twelve-bit camera (4096 grey levels) might become available to the Gallery in the future.

While it is possible to correct geometric distortions in an image using software, this process is computationally

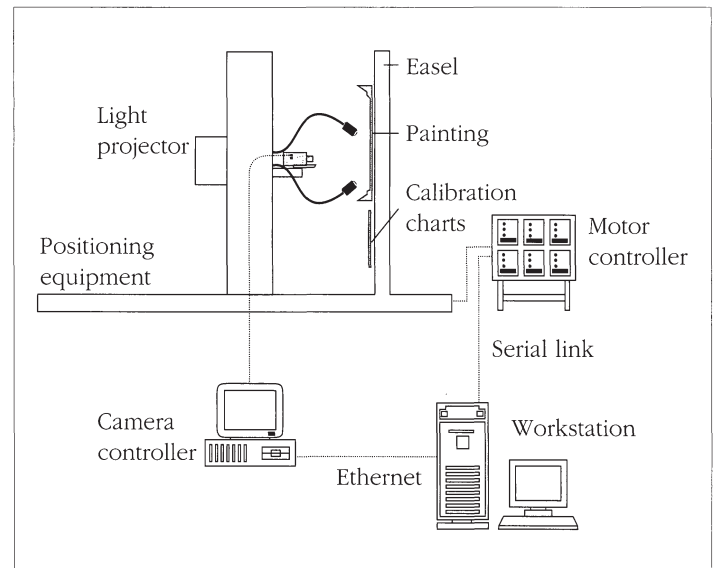


Fig. 2 Schematic representation of the image processing system.

expensive and inevitably results in some loss of image quality. A high-quality lens, which produces little or no distortion, is therefore desirable. After a search for suitable optics, a Zeiss Z-Planar f/32 lens was selected. This lens is designed to be used in instruments which produce microfiche films. It has low geometric distortion, low radiometric distortion (vignetting or shading) and a high resolution (modulation transfer function), the most important properties demanded by this application.

At present the camera control board can only operate with the bus in IBM PC-compatible computers. Images are grabbed from the camera into a seven megabyte RAM disc on a PC and transferred to the workstation over the Gallery's Ethernet LAN (Local Area Network). This transfer is the rate-determining step in the acquisition process. A board which will make it possible to




Fig. 3 Schematic representation of the operation of the Piezo Aperture Displacement (PAD) device in the ProgRes 3000 camera. The white areas are the photoactive sensors which are translated by the PAD mechanism as indicated.




Fig. 4 The lighting system and camera mounted on the main axes of the positioning equipment.

connect the camera directly to the workstation should become available in the near future.

Positioning equipment

As mentioned in the previous section, 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 the Gallery's specification by Time and Precision Ltd. It consists of a rigid steel base mounted on a heavy concrete floor with a set of vibration damping blocks. On the base are two 2.5m stainless steel rails on which the main portal moves (Fig. 1). 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 (Fig. 4). Both axes are motorised and can be moved under computer control by up to 1m in the horizontal and vertical directions to a resolution of $5\mu\text{m}$. At the far end of the base is the easel upon which paintings are mounted (see Figs. 1 and 2).²

The camera platform contains a third motorised axis which allows the camera to be moved perpendicular to the painting over a distance of about 10cm. This axis is used for fine, automatic, focusing.

The motors in the positioning equipment are driven

by a microprocessor-controlled interface, which is in turn controlled across a serial link by the workstation.³ The motor interface has two spare output lines; these are used by the workstation to control the light projector, see below.

The accuracy of the positioning equipment was assessed by TÜV-Bayern, a German state standards organisation and collaborator in the VASARI project. The repositioning accuracy was found to be better than $10\mu\text{m}$, with an absolute positioning accuracy of about $30\mu\text{m}$. 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 precisely perpendicular to the plane of the easel, and that the movement of both the horizontal and vertical axes is exactly parallel to this same 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.

Light projector

The light projector is mounted at the rear of the camera platform (Fig. 4). Because the light projector and the fibre-optic light guides which illuminate the painting move with the camera, the light distribution is identical no matter where on the surface of the painting the camera is directed. Thus, the same light-distribution correction can be applied to each sub-image. The development of the current light projector proceeded through a number of prototypes.

The first light projector contained a 12V, 100W tungsten halogen lamp in a custom-built housing. Light from the lamp was focused through an infra-red-blocking filter into the end of a fibre-optic light guide which terminated in six lenses aimed at the surface to be imaged. The filters used to produce colour separation images were placed between the painting and the camera lens (Fig. 5). The selection of suitable filters is discussed in the section on colour measurement.

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, resulting in a displacement between colour separation images of the order of several pixels. In addition, the low transmittance of these

filters led to small signals and hence high noise in the extreme red (700nm) and blue (400nm) portions of the visible spectrum.

The solution was to place the interference filters in the optical path between the light source and the fibre-optic guide. This reduced the light level to which the painting was exposed 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.

At the same time the light levels were increased by replacing the 12V, 100W bulb by a 24V, 250W slide-projector lamp fitted with a dichroic reflector. The power for the light projector and camera was drawn from a voltage-stabilised supply, after the discovery that the mains supply to the laboratory showed significant fluctuations. Additional fans were incorporated into the light projector and more efficient infra-red filters added between the source and the interference filters (Fig. 6).

Despite these modifications, the extra heat generated by the new light source created an irritating problem. It became clear that the interference filters were very sensitive to changes in temperature: in the hour or so needed to image a painting, the heat from the light source dramatically altered their transmittance characteristics.

An improved lighting system was commissioned; Fig. 7 shows the present light projector schematically. Two 24V, 250W DC tungsten halogen lamps are contained in separate enclosures. Each is equipped with a set of collimating optics and an infra-red filter. The lamps can be switched on or off under the control of the workstation. 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 infra-red output of the light source is attenuated both by the filter in the light box and by the connecting fibre-optic guide.

In the filter box the light passes through one of seven broadband interference filters. The filters are secured inside an aluminium wheel using Teflon spacers. A fan passes a steady stream of air through an elaborate set of channels cut in the filter wheel, cooling whichever filter is in the light path. 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 neutral density filter.

The filter wheel is connected to the axle of a small stepper motor. A microswitch, which senses the position of notches cut in the edge of the wheel, allows the workstation (via the positioning system interface) to select a specific filter. The light projector is illustrated in Fig. 8.




Fig. 5 Schematic representation of the prototype light projector.

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 even illumination over the region to be imaged. The complete lighting system can be seen in Fig. 4.

Workstation

The camera, the positioning equipment and the light projector are controlled by a Sun SPARCstation 2 GS workstation which is also used to gather and process




Fig. 6 Schematic representation of the modified prototype light projector.




Fig. 7 Schematic representation of current light projector.

the images from the camera. This machine has 32MBytes of RAM, a processor rated at roughly 26 million instructions per second and is able to display so-called full colour (over 16 million possible colours) images on its 1,152 by 900 pixel screen. An extra 2GBytes of hard disc storage has been added to provide adequate working memory. A Maxtor Tahiti 1GByte optical disc drive is used to archive both the raw camera data and complete calibrated images.

The software used by the workstation to acquire, calibrate and assemble images is detailed in the section on scanning procedure.




Fig. 8 The light projector. To the right are the two enclosures containing the light sources. To the left is the motorised filter wheel and fibre-optic light guide.

Colour measurement

One aim in the development of the present image processing system was to record colour information for a whole painting in a generally recognised colour notation.

Most colour measurement systems are designed to take a small number of spot readings with great accuracy. Spectrophotometers use a stabilised light source, a monochromator and a photomultiplier tube or solid-state sensor to determine the reflectance spectrum of a small area of colour. Standard colour coordinates can be derived from the spectrum of the sample. In an image processing system the requirements are very different: it is necessary to measure the colour of every 0.01 mm² in a painting of up to 1 m². A radically different approach must be adopted.

First, it is necessary to modify the response of the Kontron ProgRes 3000 monochrome camera with filters in order to obtain a series of colour separation images. Second, the colour separation images must be processed by the workstation to yield colour coordinates in the chosen notation.

Spectral reconstruction

It is theoretically possible to record the reflectance spectrum for each pixel in the image using the following relationship between the response of the camera and the reflectance of the object colour:

$$\chi_k = \Phi \left[\int_{\lambda_{\min}}^{\lambda_{\max}} \rho(\lambda) \cdot s(\lambda) \cdot \phi_k(\lambda) \cdot o(\lambda) \cdot c(\lambda) \cdot d\lambda \right]$$

- χ_k = camera response through filter k
- $\rho(\lambda)$ = reflectance of object colour at wavelength λ
- $s(\lambda)$ = power of light source at wavelength λ
- $\phi_k(\lambda)$ = transmittance of filter k at wavelength λ
- $o(\lambda)$ = transmittance of camera optics at wavelength λ
- $c(\lambda)$ = sensitivity of photodetector at wavelength λ
- Φ = non-linearity correction function

A derivation of the complete solution of this equation for the visible region (400 to 750nm) with a resolution of 10nm, which requires a set of 36 filters and 36 reference colours whose reflectance spectrum is known, is outside the scope of this present article. A series of studies conducted by the Ecole Nationale Supérieure des Télécommunications (ENST) in Paris, a collaborator in the VASARI project, indicated that the minimum number of filters for spectral reconstruction from colour separation images was twelve.⁴ It was also shown that since artists' pigments have very




Fig. 9 Transmittance curves for the seven broad band interference filters used in the light projector.

smooth spectral curves, only a few points are necessary to characterise their spectra with good accuracy.

A further theoretical study was conducted at the National Gallery to compare the set of twelve filters recommended by ENST with a set of narrow band interference filters, a set of broad band interference filters and a set of three traditional red, green and blue colour separation filters.⁵ The 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. It can be seen from Fig.9 that their transmittances cover the visible spectrum with considerable overlap. Hence, most of the spectral information will be contained in the seven-channel image recorded using this filter set.

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. The time taken to record twelve images would be proportionally longer, a larger filter wheel would be required to accommodate the increased number of filters and finally the data generated by recording twelve colour separation images for a painting at a resolution of 10 pixels per millimetre would be approximately 1.2GBytes per square metre.

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 recognised colour notation. The colour system chosen is the Commission Internationale de l'Eclairage (CIE) XYZ standard colour space.⁶ Since other CIE colour coordinates, for instance $L^*a^*b^*$, may be derived from the values of XYZ, the latter seem an obvious choice.

Because the process of reconstructing the reflectance spectrum prior to calculating the colour coordinates is extremely computationally intensive, it was decided to adopt a different approach and to derive XYZ (the so-called tristimulus values) directly from the response of the camera through each of the seven filters.

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:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = F \cdot \begin{pmatrix} \chi_1 \\ \chi_2 \\ \vdots \\ \chi_7 \end{pmatrix}$$

where

$$F = \begin{pmatrix} f_{X1} & f_{X2} & \dots & f_{X7} \\ f_{Y1} & f_{Y2} & \dots & f_{Y7} \\ f_{Z1} & f_{Z2} & \dots & f_{Z7} \end{pmatrix}$$

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 ($\chi_{11}, \chi_{12}, \dots, \chi_{7n}$) 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. Thus:

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

where M^T denotes the transpose of matrix M and

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

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.

Colour standards

The selection of the colour standards used to derive the conversion matrix has received considerable attention. Initially a set of ceramic colour standards from the National Physical Laboratory (NPL) was used.

Although these tiles have a variety of hues, saturations and lightnesses, they do not cover the range of colours found in easel paintings. To supplement these, additional tiles with a wider range of hues were acquired. Unfortunately, unlike the NPL tiles which are designed to have a homogeneous and permanent colour, their surface colour was not uniform.

The present colour reference used is a Macbeth ColorChecker Chart®. This chart comprises 24 colour patches, including the additive and subtractive primaries and a grey scale. The stability and homogeneity of the colours are 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.

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).⁷ These may be derived from the values of XYZ stored for each pixel.⁸ 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'.⁹

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 British Standard 6923:1988, attempts to compensate for the non-uniformity of the CIELAB space.¹⁰ The derivation of the recommended formula is given in Appendix 2. It is sufficient to note here that using the CMC colour difference (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 throughout the VASARI project.

Accuracy of colour measurement

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 sets of 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.

Scanning procedure

A colorimetric image of a painting is built in two distinct phases. In the first, *acquisition*, raw data is taken from the camera and stored on the workstation's hard disc. In the second phase, *calibration*, the workstation analyses the data produced during acquisition. 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 is a set of calibration targets. These are shown in Fig. 10 and are:

- 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, as described above, which is used to calibrate the colour in each sub-image of the painting.
- A grey level reference; a standard Kodak grey scale, which is used to detect non-linearities in the response of the camera's CCD sensor.

Acquisition

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 (Fig. 10).

- 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 2.5 m 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 at a series of points 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. ACQUIRE executes the following 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. This is equivalent to choosing an aperture setting, but is more controllable.
- 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. These images are used in the calibration phase to correct non-linearities in the response of the CCD sensor.




Fig. 10 The easel. At the top left is the white target, beneath which is the resolution target. At the base of the easel on the left is the colour calibration target. To its right is the grey scale. The painting to be scanned is held to the right of the calibration targets.

- 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 acquired through each of the seven filters.

Correction and calibration

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.

- 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.
- Finally, each sub-image is converted to XYZ colour space using the colour conversion matrix.

Mosaic assembly

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 (Fig. 11a). From this information the program calculates the region of overlap between the two sub-images (Fig. 11b). 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 (Fig. 11c). 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 (Fig. 11d). In the region of overlap the computer merges data from the two images to give a smooth transition (see Plate 1, p. 82).

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.

Display

The XYZ image that has been generated by the calibration procedure cannot itself be displayed directly. The electron gun voltages for the red, green and blue phosphors in the CRT (cathode ray tube) monitor must be derived from the XYZ values in order for the correct colours to appear on the screen. To compute this transformation, it is necessary to know

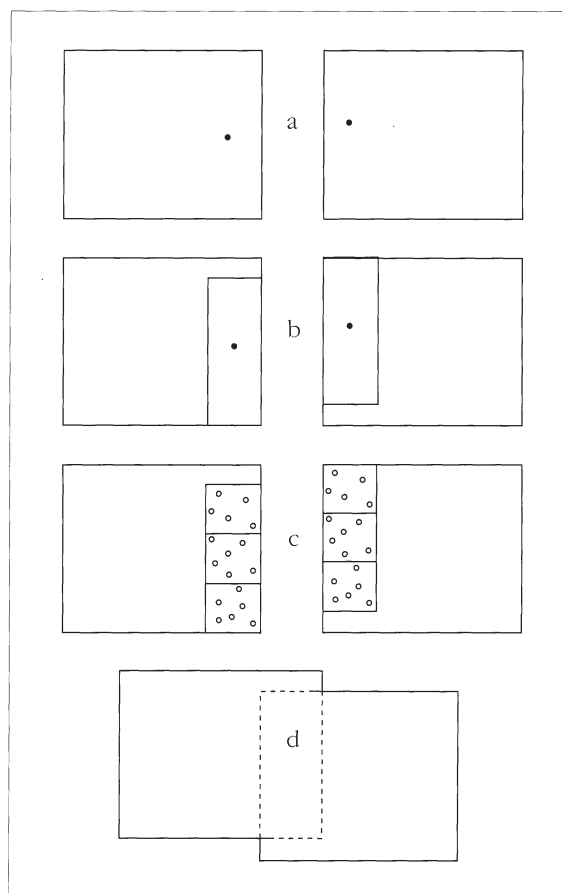


Fig. 11 The image assembly procedure.

- a) Corresponding points in two adjacent images are selected by the operator.
- b) The region of overlap is calculated automatically.
- c) Further tie points are selected by the computer.
- d) The two images are merged.

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 transformation from XYZ to *rgb* (red, green and blue gun voltages) calculated from this information will not be detailed here.¹¹

There is an additional problem: the gamut of colours that can be generated by a CRT monitor is much smaller than the gamut of colours found in paintings.¹² 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 gamut of the screen. A full description of this 'fitting' process is available elsewhere.¹³

The large size of the final mosaic often 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 stored on the hard disc (see Plate 2, p. 82).

At present, the Gallery does not possess the equipment needed to generate 'hard copies' of images. Black and white negatives or colour transparencies can be produced by sending the digital data (on tape or disc) to an outside reprographic company but it is very difficult to control the quality of such reproductions.

Storage

The images are stored temporarily on the computer's hard disc, prior to permanent archiving on optical disc. The raw data for each of the seven channels are archived as well as an XYZ and *rgb* record of the whole painting. The low-resolution version of the *rgb* image is also stored to provide a quick reference.

Applications

The VASARI system is being used for a number of research programmes. In order to facilitate this work, a package of image processing programs has been written which allows non-specialist users to manipulate images on the workstation.¹⁴ Plate 3 (p. 83) illustrates the use of one of the image processing routines.

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 infra-red camera to the system, it has been possible to make high-quality infra-red reflectogram images of the underdrawing in paintings. These and other applications are described in more detail in the sections which follow.

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, may 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 program.

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. The problem can be ameliorated by maintaining the climatic conditions in the area in which scanning takes place, but this 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 'yellower' than the corresponding pixels in the old image might be shown by producing a map showing the increase in CIE b^* value.

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

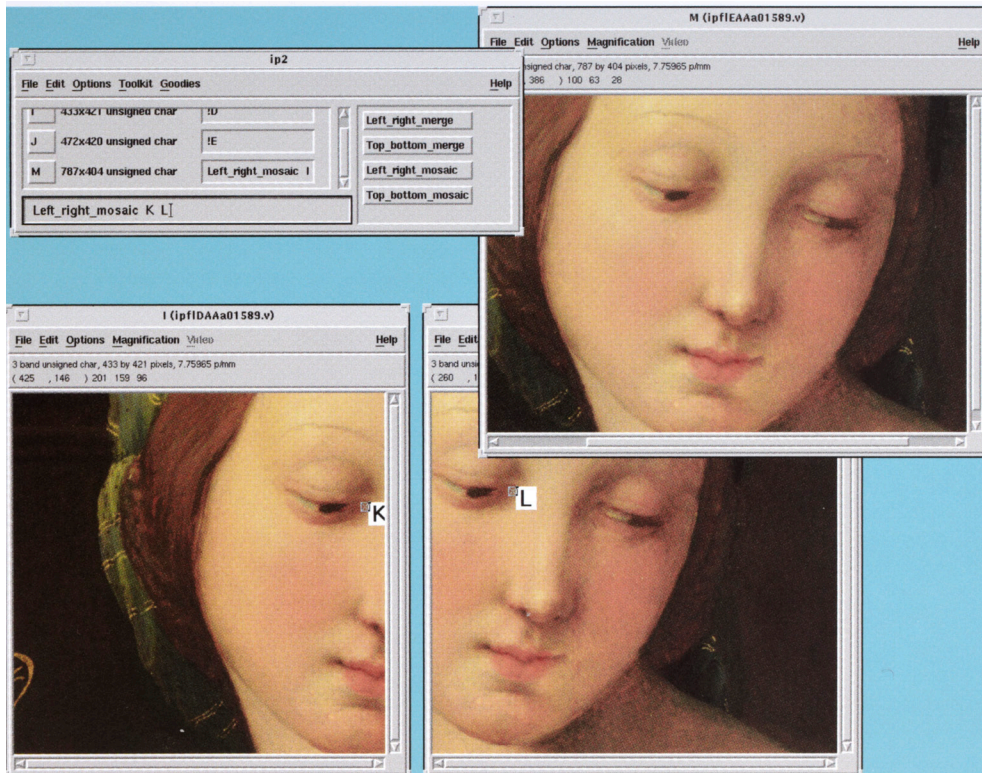


Plate 1 The mosaic assembly procedure. To the bottom of the screen are two sub-images. At the top of the screen is the merged image. Note the points selected by the operator in each sub-image, which have been used in the mosaic assembly program.



Plate 2 Selecting a detail. The low-resolution image at the upper right of the screen is used to select a detail from the high-resolution image. The remainder of the screen is occupied by the portion of the high-resolution image thus selected.

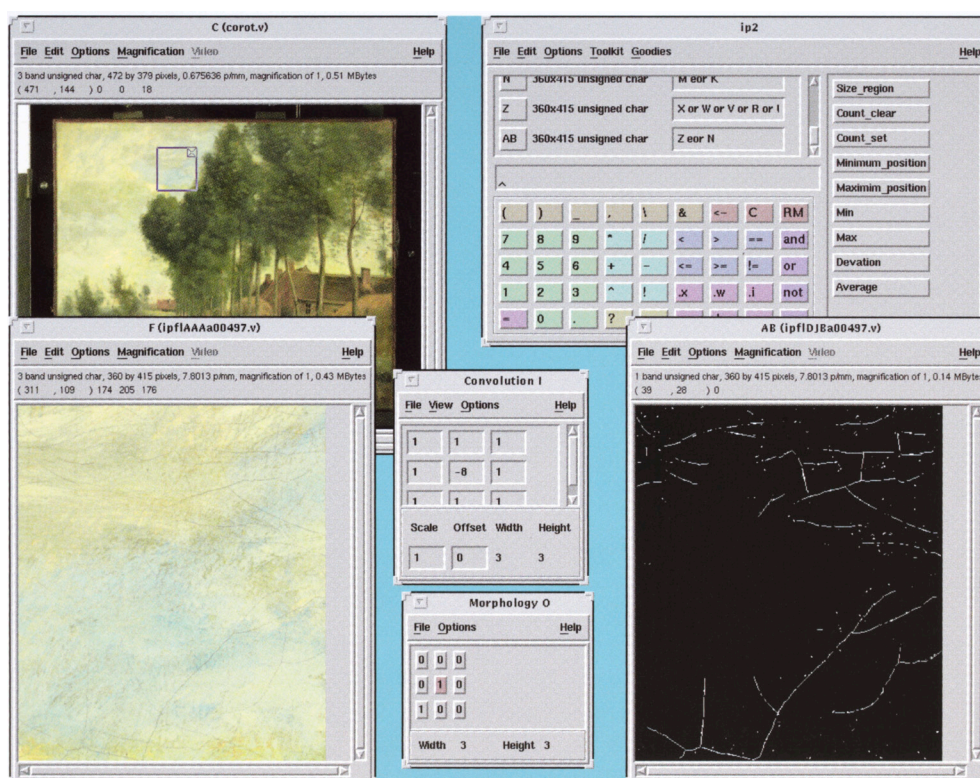


Plate 3 An example of the use of the image processing package. A detail from a high-resolution image is displayed to the left of the screen. The operator has selected a craquelure extraction procedure using the dialogue boxes alongside. The processed image, showing the crack pattern, is to the right of the screen.

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 road.¹⁵

Infra-red reflectography

The image acquisition and processing software has also been used in the examination of paintings by infra-red reflectography. An infra-red vidicon camera has been attached to a second workstation via a digitising board. The infra-red sub-images can be corrected and assembled into a mosaic more rapidly and with greater precision than is possible using the conventional photographic method. The procedure is described in more detail elsewhere.¹⁶ It is being used in the 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-images. The infra-red reflectogram of Raphael's '*Garvagh Madonna*' (see p.12 of this *Bulletin*) was computer-assembled in this way.

Digitisation of photographic material

By placing a light-box on the easel it is possible to use the VASARI equipment to digitise the negatives of existing photographs or radiographs. New X-radiograph mosaics have been assembled. In addition it has been possible to use the image processing software to analyse the periodicity of patterns in X-radiographs of canvases and in β -radiographs of paper texture.¹⁷

Future developments

The National Gallery's image acquisition and processing system is being used to begin a systematic recording of paintings from the Collection. Planned modifications include the installation of a twelve-bit A-D converter board and an increase in the hard disc capacity of the workstation. As computer technology advances it is to be expected that larger capacity optical storage devices will become available.

Some improvements in the procedure for calibration are envisaged. A new, more efficient and rapid version of the calibration software is nearing completion. Software to automate the registration and comparison of images will be developed to expedite the analysis of change in paintings. One drawback with the computer

system at the National Gallery is that it is not widely available in other museums and galleries. In the future it might be possible to create versions of the software for the more conventional Macintosh or IBM PC-compatible computers.

At present colour images stored on the system cannot easily be reproduced on paper with good colour fidelity. The successor project to VASARI, entitled MARC (Methodology for Art Reproduction in Colour), is investigating the accurate reproduction of colorimetric images on paper. These results will be of use to the Gallery as it considers establishing in-house desktop publishing facilities for books and catalogues.

Acknowledgements

Although the majority of the work described in this paper has been conducted at the National Gallery, it results from the combined effort of a number of the collaborators in the VASARI project. In particular, Drs Kirk Martinez and Nicolaos Dessipris of Birkbeck College, University of London, were responsible for much of the general image processing software; Dr Andreas Burmester and Manfred Müller of the Doerner Institut, Munich, developed software for craquelure detection and image alignment; Professor Henri Maitre and Dr Francis Schmidt of ENST conducted theoretical studies into colorimetric imaging methods.

Appendix 1 : The VASARI consortium

This paper presents some of the work of the VASARI consortium carried out as part of ESPRIT II project number 2649. Other institutions involved in this project were: Birkbeck College, University of London (UK); The Doerner Institut, Munich (D); Brameur Ltd. (UK); ENST [Telecom Paris] (F); Thomson-CSF, Rennes (F); TÜV-Bayern, Munich (D).

Appendix 2 : Derivation of the CMC formula

The hue angle, h_{ab} , and chroma, C_{ab}^* , are calculated from a^* and b^* as follows:

$$h_{ab} = \arctan (b^*/a^*)$$

$$C_{ab}^* = (a^{*2} + b^{*2})^{1/2}$$

The hue difference between the two colours is determined from the CIELAB colour difference ΔE_{ab}^* ,

$$\Delta H_{ab}^* = (\Delta E_{ab}^{*2} - \Delta L^{*2} - \Delta C_{ab}^{*2})^{1/2}$$

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

$$\Delta E_{CMC}^* = ((\Delta L^*/S_L)^2 + (\Delta C_{ab}^*/S_C)^2 + (\Delta H_{ab}^*/S_H)^2)^{1/2}$$

where $S_L = 0.040957L^* / (1 + 0.01765L^*)$

unless $L^* < 16$ in which case $S_L = 0.511$

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

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

where $f = (C_{ab}^{*4} / (C_{ab}^{*4} + 1900))^{1/2}$

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

unless h_{ab} is between 164° and 345° in which case

$$T = 0.56 + |0.2 \cos(h_{ab} + 168)|$$

Notes and references

1. D. Saunders, 'Colour Change Measurement by Digital Image Processing', *National Gallery Technical Bulletin*, 12, 1988, pp.66–77.
2. The positioning system installed at the Doerner Institut in Munich, a collaborator in the VASARI project, can scan paintings of up to 1.5 m by 1.5 m.
3. All the computer programs described in this paper run on the workstation under SunOS Unix 4.1 (Sun operating system) or greater, are written in the C computer language and the graphical tools use the Motif 1.1 user interface toolkit under X11R4 window system. A program entitled XMCONTROL provides a mouse-based interface to the camera and the positioning equipment. The user can move the camera over the painting with the mouse, change filters, turn the lights on and off and grab images at varying resolutions. Other applications running on the workstation, such as ACQUIRE and CALIBRATE, use the facilities provided by XMCONTROL to operate the positioning equipment, camera and light projector under computer control.
4. G. Deconinck, 'Multi-channel Images: Acquisition and Coding Problems', MSc thesis, Katholieke Universiteit Leuven and Ecole Nationale Supérieure de Télécommunications, 1990.
5. D. Saunders and A. Hamber, 'From Pigments to Pixels: Measurement and Display of the Colour Gamut of Paintings', *Proceedings of the Society of Photo-Optical Instrumentation Engineers*, 1250, 1990, pp.90–102.
6. 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.

7. See Note 6 above.

8. The CIE L*a*b* colour coordinates are derived from XYZ as follows:

$$L^* = 116(Y / Y_n)^{1/3} - 16 \quad \{\text{for } Y / Y_n > 0.008856\}$$

$$L^* = 903.3(Y / Y_n) \quad \{\text{for } Y / Y_n \leq 0.008856\}$$

$$a^* = 500[(X / X_n)^{1/3} - (Y / Y_n)^{1/3}]$$

$$b^* = 200[(Y / Y_n)^{1/3} - (Z / Z_n)^{1/3}]$$

Where X_n , Y_n and Z_n are constants which depend upon the light source for which the calculation is being made.

9. See Note 6 above.

10. F. J. J. Clarke, R. McDonald and B. Rigg, 'Modification to the JCP79 Colour-Difference Formula', *Journal of the Society of Dyers and Colourists*, 100, 1984, pp. 128–32.

11. The details are contained in D. Saunders and J. Cupitt, 'Colour Display Techniques: Part II', *VASARI Report R-CCDM1.2*, 1990, pp. 1–4. See also D. L. Post and C. S. Calhoun, 'An Evaluation of Methods for Producing Desired Colors on CRT Monitors', *Color Research and Applications*, 14, 1989, pp. 172–86.

12. See Note 5 above.

13. D. Saunders and J. Cupitt, op. cit., pp. 5–14.

14. Since no easily extensible commercial system

could be found which was capable of processing images of the necessary size, an image processing package designed to be used with images of works of art was developed during the project. The package is entitled *VIPS* (VASARI Image Processing System). *VIPS* provides the usual range of operations found in image processing packages, including a comprehensive range of arithmetic operations, morphological operators, colour change operators and filtering operations. The basic image processing routines in *VIPS* are used by *ACQUIRE* and *CALIBRATE* for image handling. A convenient, graphical user interface for general image processing has been developed by Thomson-CSF, a partner in the VASARI consortium. Operations are executed by selecting, with the mouse, an icon representing the action to be performed. A second interface, called *IP*, has been developed at the Gallery for scientific and technical users.

15. A description of this preliminary study is contained in A. Burmester and M. Müller, 'The Registration of Transportation Damages using Digital Image processing', *Zeitschrift für Kunsttechnologie und Konservierung*, 6, 1992, pp. 335–45.

16. 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*, accepted for publication.

17. A report on the application of image processing techniques to β -radiographs of the paper in Van Dyck's Antwerp sketchbook is in preparation.