

# Ten Years of Art Imaging Research

KIRK MARTINEZ, MEMBER, IEEE, JOHN CUPITT, DAVID SAUNDERS, AND RUVEN PILLAY

## Invited Paper

*This paper describes a decade of work on digital imaging for museums. From 1989 to 1992, the visual arts system for archiving and retrieval of images (VASARI) project produced a digital-imaging system that made color-calibrated images of up to  $20\,000 \times 20\,000$  pixels directly from paintings. It used seven color-separation bands in the visible region, resulting in an average color error of around  $1 \Delta E_{ab}^*$  unit. These images have since been used to monitor the condition of paintings, document paintings during conservation treatment, including predicting appearance after cleaning, reconstruct the original appearance of paintings in which pigments have faded, and assess whether paintings have been damaged during transportation, in estimations of the surface reflectance spectra and in the printing of high-quality reproductions. We have applied similar techniques to museum infrared and X-ray imaging. To manage the images produced by the VASARI system, an image-processing package has been developed that is tailored for very large colorimetric images. This package has since been used in several other projects, including a remote image viewer designed to provide internet access to high-resolution images. The paper explores these developments and gives details of the current generation of VASARI-derived systems, set in the context of the state of the art for museum imaging.*

**Keywords**—Colorimetry, image processing, imaging, museum, painting.

## I. INTRODUCTION

This paper summarizes roughly a decade of work developing imaging systems for digitizing works of art, particularly paintings. Although these systems were built specifically to image paintings in museums and galleries for scientific research purposes, the techniques used are more widely applicable having found use in the documentation

of works of art and for the dissemination of high quality images. The expertise needed for their realization was drawn from a number of fields, including optics, mechanical engineering, image processing, color science, computer science, and painting conservation.

A principal motivation behind the original project was to provide an accurate means of measuring the color across the entire surface of the painting, to provide a definitive record of the state of an object at a given time. Against this image, future images recorded with equal precision could be compared to give an indication of any changes that had occurred. When we started the research, the only accurate way to measure color accurately on paintings was to take spot readings of spectra using a spectrometer and record a rough position of the sample on a photograph, using a technique that had been developed at the National Gallery, London, U.K., and Courtauld Institute of Art, London, U.K., by Wright and Wassal [1]–[3]. We expected the color accuracy of our new system to be lower than spectrophotometry, but felt that the ability to detect color change anywhere on the surface was more important. With spot techniques, it is only possible to measure the color at a limited number of points on the surface of the painting. This means that measurement sites have to be selected very carefully: one has to decide at the outset which areas of the painting are expected to change color during the next 20 years—not an easy problem. Imaging techniques avoid this difficulty by simply measuring the entire surface. Indeed, for this reason, two prototype image-processing systems had already been built and tested at the National Gallery during the 1980s [4], [5]. Although the data provided by these systems was self-consistent, neither produced color measurements in a standard color space that could be compared with colorimetric measurements made using other techniques.

By the end of the 1980s, digital imaging was developing at a rapid pace. New devices such as charge-coupled devices (CCDs) made it possible to consider replacing not only spectrophotometric measurement of color by direct digital imaging, but also to consider whether digital images might one day replace conventional photographs in the museum. Digital imaging has compelling conservation advantages over film. Although the usual records of paintings are large

Manuscript received July 12, 2000; revised July 5, 2001. This work was supported in part by the European Commission.

K. Martinez is with the Department of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, U.K. (e-mail: km@ecs.soton.ac.uk).

J. Cupitt and D. Saunders are with the National Gallery, London WC2N 5DN, U.K. (e-mail: john.cupitt@ng-london.org.uk; David.Saunders@ng-london.org.uk).

R. Pillay was with The National Gallery, London WC2N 5DN, U.K. He is now with the Central Research Laboratories, the Louvre, 75058 Paris Cedex 01, France.

Publisher Item Identifier S 0018-9219(02)00733-8.

format transparencies (up to 10 in  $\times$  8 in), which have excellent resolution, they gradually deteriorate over time and have inaccurate color. This usually means that objects need photographing regularly to replace old photographs, a process that may expose the painting to more light and handling. Digital images promised permanence and color accuracy ensuring that images taken many years apart could be compared. The easy transmission of digital images make sharing information with publishers and conservators in partner institutions far easier.

Conventional records of paintings consist of many different types of photograph. For example, in the laboratory of the Louvre Museum in Paris, dossiers for each painting contain photographs made under ultraviolet, visible, infrared, and X-ray radiation. These images come in many different formats: infrared images are often offscreen photographs, while X-radiographs are typically very large format negatives similar to those used in medical applications. These films are extremely valuable and may be the only image left if a painting is destroyed or damaged. Large-format color transparencies are used by museum and galleries as the basis for their printing and multimedia projects, although it is not uncommon for a new transparency to be made of a detail, for example, for a specific project.

Digital images offer the possibility of storing high-quality images that do not need to be scanned prior to reproduction, a process that inevitably leads to the deterioration of the transparency. It also makes comparisons between images more straightforward. To replace 10-in  $\times$  8-in transparencies with digital imaging was a daunting task. In the National Gallery, for example, these are taken in a studio, processed, assessed, and retaken if not acceptable; they are, therefore, of very high quality. We believed, however, that digital imaging could give improvements on conventional photography, for example, by removing lighting variations and standardizing the color in different images.

At the end of the 1980s, the main challenge that we faced was to construct a digital imaging solution that had adequate spatial resolution to allow us to make long term studies of color change and that also offered a credible alternative to high quality conventional photographs. From the point of view of a photographer, to replace a standard 10-in  $\times$  8-in transparency film resolving 40 lines pairs per millimeter would require an image of around 16 000  $\times$  20 000 pixels. From the point of view of a conservator, the smallest feature of a painting that you might reasonably want to record is the pattern of cracks (cracquelure) in the surface layers of paint. Examinations of X-radiographs showed that cracks measuring around 0.1 mm are typical. Therefore, a resolution of 20 pixel/mm on the surface of the painting is required. For a 1-m  $\times$  1-m painting, 20 pixel/mm corresponds to 20 000  $\times$  20 000 pixels. This was far higher than any imaging device available in 1989.

Another difficult question was the format to be used for storing color information. We needed to be certain that our measurements would be useful in five or even 20 years time, when our successors would need to compare our readings against their own, almost certainly made with very different

equipment. The ideal would be to store reflectance spectra, but to allow straightforward color comparisons to be made using a standard system, all the color data have been converted into the CIE Lab [6] color space. Using the results of vision research, one accuracy target was to obtain an average color error of less than  $\Delta E_{ab}^*$  so that color errors would not be discernible to the human eye. If the ideal of spectral data was achieved, then new capabilities such as showing the image as it would appear under different light would become possible.

#### A. VASARI Project

In 1989, we were part of a group that was interested in exploring the limits of digital imaging. With colleagues from other European museums and universities we started a European Community-supported project called the visual arts system for archiving and retrieval of images (VASARI) [7]–[9]. The project aims were to produce an imaging system good enough to replace technical photography of paintings and to use this system to look for changes in paintings on display or changes occurring during transportation between museums [10]. The side benefits of having a stock of high-resolution digital images would be easy access to details without rephotography, for printing and CD-ROM production, for example. With a high-resolution image, a detail can simply be extracted from the file instead of having to photograph the painting again. The goals of VASARI were sufficiently difficult to inspire either incredulity or amazement in other people, so we were lucky to be funded. It has to be remembered that we were talking about using the capacity of one CD-ROM to store one image, with hard disks still at the 1-GB stage, so many people were concerned about the storage requirements. The 2000 paintings in the National Gallery would occupy around 1 TB and museums such as the Louvre have more than ten times that number.

The project resulted in the production of two imaging systems: one installed in the National Gallery, optimized for color measurement and the second in the Doerner Institute, Munich, Germany, optimized for monochrome transportation damage assessment. As a result of a further European initiative (MUSA European Project 7523), a third scanner similar to that in London was installed at the Galleria degli Uffizi, Florence, Italy.

1) *Hardware Design:* At an early stage of the design, we decided that the painting should be scanned while vertical, since the shape of a canvas painting changes when it is laid flat. This was in contrast to comparable systems, such as the innovative scanner that had already been developed to digitize the U.S. Constitution in Washington, DC. [11]. As well as maintaining the paintings in their normal position, this would also help to prevent accidental damage if sections of the prototype scanner were to drop off.

We hoped to be able to reproduce spectra so we went ahead with a multispectral approach rather than the conventional red-green-blue (RGB) color separation. An ideal RGB set is impossible to achieve and in theory the multispectral technique can provide better colorimetric values. Our partners at Ecole Nationale Supérieure des Télécommunications




Fig. 1. Schematic of the VASARI scanner lighting system.

(ENST), Paris, carried out simulations using the spectra of known artists' pigments (provided by the National Gallery) to define an ideal set of filters for imaging paintings. This filter set was optimized to allow the reconstruction of the (relatively smooth) spectra of traditional artists' pigments [12]. Unfortunately, these 12 theoretical filters had no affordable or readily available real counterparts. A pragmatic solution was, therefore, adopted, which gave good coverage of the visible spectrum and used filters that transmitted sufficient light to provide an acceptable signal from the camera without overexposing the painting to light. We selected seven broadband Gaussian filters with peak transmittances at 400–700 nm in 50-nm intervals and a bandwidth at half maximum transmittance of 70 nm. Simulations showed that this Gaussian filter set was only slightly worse than the ideal set. We were also confident that the system could be upgraded at a later date, as better filters lights or optical fibers became available (all of which have been replaced in the course of ten years use). Although, in theory, more accuracy can be achieved with more filters [13], a practical limit is imposed, as described above, by the amount of light required to provide a signal in relation to the exposure of the painting to potentially damaging radiation.

The interference filters we chose are sensitive to heat, making it impossible to get enough light through a single filter to illuminate an entire painting. The filters we chose also were rather thick and varied greatly in thickness, making them very hard to use inside the optical path of a camera. Our solution was to use fiber-optic guides to pass light through a single filter and then illuminate a small patch of the painting. This had the additional important advantage of exposing the painting to far less light during scanning. Initial experiments with the filter close to the light were problematic, as the filters tended to crack with the heat. Even under conditions where the filters were not permanently

damaged, the heating of individual filters as they passed into the light path resulted in temporal changes in the spectral characteristics of the transmitted light. An infrared blocking filter was placed in front of the light source and fiber-optic light guides were used to transmit “white” light from the projectors to an enclosure containing the filter wheel. In addition, a fan drew cool air across the filter while it sat in the light path to minimize the effect of local heating. A light source with a continuous spectrum was selected; we used easily available tungsten-halogen lamps. Main supply stabilization was also necessary due to the long exposure times and nearby restaurant causing spikes. Fig. 1 shows a schematic of the lighting system.

In order to obtain high resolution, we were going to image small portions at a time anyway, but it turned out that due to the lighting system we would only have enough light to illuminate a small part of the painting at any time anyway. One obvious design idea was to use a stationary camera and lighting system and to move the painting in front of the camera on a computer-controlled easel. Again, for safety reasons, we felt that we had to keep the painting stationary. This meant that we had to transport the entire camera and lighting system (about 15 kg) in front of the surface of the painting. We were unsure what accuracy we would require for this motion, so when we commissioned the positioning system, we asked for 50  $\mu\text{m}$  based on an estimate of the compensation we would be able to make in the software stitching and the resolution for which we were aiming. In retrospect, the positioning accuracy of the scanner could have been considerably lower, as the software, described later, is able to compensate for positional inaccuracies. The most important feature of the system turned out to be that it was sufficiently rigid to maintain coplanarity between the painting being scanned and the focal plane of the camera. Although considerable setting up is required to make sure

that this condition is fulfilled, the alternative (trying to make mosaics from trapezoid images) is extremely difficult.

We selected the highest resolution camera we could afford, the Kontron ProgRes 3000 [14]. This camera could take a  $3000 \times 2000$  pixel color image in about 30 s, fast for the time. We bought a monochrome version and selected a lens with almost no geometric distortion and low chromatic aberration. Stitching software would be written to join the image tiles and we wanted to avoid complex geometric corrections. We had to compromise with our requirements for 12-bit analog-to-digital (A/D) conversion in the camera and use only 8 bits.

2) *Software Design:* All the software was developed in C under Unix on Sun workstations. We implemented an automatic acquisition and calibration package for the VASARI scanner. The painting to be scanned was placed on the easel, its dimensions were entered, any adjustments to the acquisition parameters were made, and scanning started. Because the camera axis ( $Z$ ) could also be moved by 10 cm, it was used to reset the focus on the painting automatically. This used a simple sum of a high-pass filter in the center of the raw low-resolution camera output to give a measure of sharpness. Measurements taken at the extremes of position were then interpolated and used in an iterative process to narrow down the best position. This meant the lens could be left in one position and there was more consistency in positioning and hence resolution.

a) *Calibration targets:* A number of calibration targets were scanned before acquisition proper started. First, the system prompted the user to cover the camera lens before taking a dark current image. This was subtracted (with an offset to prevent clipping of negative noise) from all subsequent images.

The resolution of the image depended on the camera-to-object distance, which varied slightly between acquisitions. The first target was a vertical white line on a black background. The camera moved left and right a small amount and measured the movement of the image of the line on its sensor. Scanning resolution was calculated from this. This resolution figure was used to guide the movement of the camera over the painting, to ensure that there was a small overlap between neighboring images. Images would have their resolution field set to this measurement, which was invaluable afterwards as the real scale was always available.

The fiber-optic lighting system did not produce perfectly uniform illumination so the next calibration target was a polished sheet of PTFE (nonstick plastic). This had to be perfectly clean because any marks on it would be transferred to the images. An image of the PTFE sheet was used to correct for variations in pixel sensitivity, illumination homogeneity, and lens shading in all subsequent images. Although this could have been done by a separate characterization of the CCD and lens using an integrating sphere for example, it was a one-step solution to the three problems.

A Macbeth ColorChecker chart [15] (see Fig. 2) was used for color correction[11]. Although this chart has only 24 patches, the patches are made from painted pigments rather than being produced photographically or through printing.



Fig. 2. Macbeth ColorChecker chart.




Fig. 3. VASARI scanner as it appeared in early 1992 when the hardware design was complete.

We had established that the spectral characteristics of the patches on the ColorChecker chart were not dissimilar to those of artists' pigments or paintings. We did experiment with the use of color calibration charts based either on a number of painted out artists' pigments or on a large number of color chips from the NCS system [16], but found that the quality of the calibration was not significantly improved. We cut the chart tiles into quarters and constructed a small version of the chart to fit within the field of view of the camera (a small version is now available). The chart's spectrum was measured and entered into our calibration files.

b) *Painting acquisition:* The camera was then moved over the surface of the painting, capturing seven monochrome images of  $3000 \times 2000$  pixels at each position. The camera's field of view covered an area on the painting of about  $17 \text{ cm} \times 13 \text{ cm}$ , giving a resolution of approximately 18 pixel/mm. Scanning took around three hours for a  $1\text{-m} \times 1\text{-m}$  object. Fig. 3 shows the scanner shortly after construction.

3) *Processing and Storage:* After the painting had been captured, a second processing stage was run. This used the

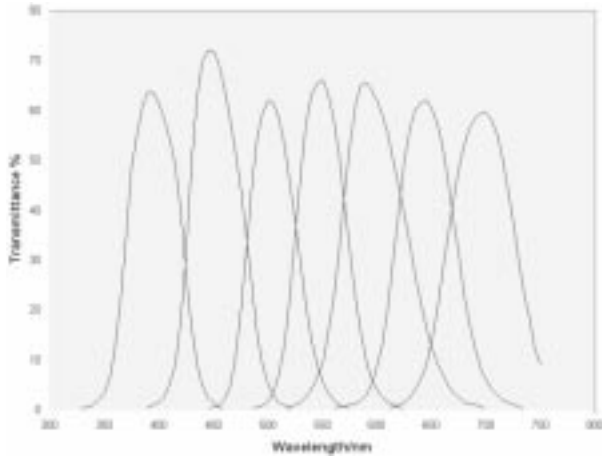


Fig. 4. Transmittance of the VASARI filters.

images captured during acquisition to automatically calculate a single very large output image. This is stored uncompressed as CIE Lab in our own simple VASARI image-processing system (VIPS) [23] file format.

First, the images of the Macbeth chart, combined with spectrophotometric measurements of the chart patches, were used to determine a transformation, which mapped the seven bands of the scan to D65 CIE  $XYZ$  color space. The camera outputs are initially linearised with respect to  $Y$ .

It is assumed that the standard observer curves used to generate CIE  $XYZ$  can be constructed with a linear combination of the Gaussian filters we used. Their transmittance can be seen in Fig. 4. Although a completely accurate reconstruction of the standard observer curves is not possible, the fact that the input spectra are so smooth leads to an accurate result in the end. A simple linear model is used where the camera values  $F$  are related to the CIE  $XYZ$  color values of the test chart  $V$  in a simple linear system where  $K$  is an unknown set of coefficients

$$V = K \cdot F$$

or in full form

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} k_{X1} & k_{X2} & \cdot & \cdot & k_{X7} \\ k_{Y1} & k_{Y2} & \cdot & \cdot & k_{Y7} \\ k_{Z1} & k_{Z2} & \cdot & \cdot & k_{Z7} \end{pmatrix} \cdot \begin{pmatrix} f_1 \\ f_2 \\ \cdot \\ \cdot \\ f_7 \end{pmatrix}.$$

The color calibration calculates the values of  $k_{X1}$  to  $k_{Z7}$  in the conversion matrix ( $K$ ) above. This is achieved by imaging  $n$  (where  $n \geq 7$ ) color standards (with known  $XYZ$  values) through each of the filters. The matrix ( $\Phi$ ) representing the camera response to each color in the chart through each filter ( $f_{11}, f_{21}, \dots, f_{7n}$ ) is combined with the known  $XYZ$  values for the  $n$  colors ( $X_1, Y_1, Z_1, \dots, X_n, Y_n, Z_n$ ), matrix ( $V$ ), to generate the least mean square (LMS) solution for the conversion matrix  $K$

$$K = [(\Phi^T \Phi)^{-1} \Phi^T V]^T$$

where

$$\Phi = \begin{pmatrix} f_{11} & f_{21} & \cdot & \cdot & f_{71} \\ f_{12} & f_{22} & \cdot & \cdot & f_{72} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ f_{1n} & f_{2n} & \cdot & \cdot & f_{7n} \end{pmatrix} \quad V = \begin{pmatrix} X_1 & Y_1 & Z_1 \\ X_2 & Y_2 & Z_2 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ X_n & Y_n & Z_n \end{pmatrix}.$$


$K$  is then used to convert images of paintings into calibrated  $XYZ$  values. The equation for  $K$  above gives an LMS solution which usually gives an average color error on the Macbeth chart of  $1 \Delta E_{ab}^*$  units. Most colors from the chart are calibrated very well but others may be as far as  $5 \Delta E_{ab}^*$  units in error.

Next, each of the image tiles making up the painting were transformed to CIE Lab and saved back to disc. Finally, the overlapping areas between tiles were analyzed and the correlation surfaces used to find a set of joins that minimized errors between neighboring subimages [17]. The image patches are mosaiced together using normalized correlation of the overlapping areas. In order to ensure that the correlation uses image areas that are not smooth, a Laplacian high-pass filter is applied to the overlaps and around 20 areas of high detail found. These smaller areas are then correlated with slightly larger (by 20 pixels, typically) corresponding areas on the other image. These search-area limiting techniques, together with knowledge of the overlap size (typically  $100 \pm 20$  pixels), increases the success rate. By using three main regions for correlation (top, center, and bottom), correlated features provide points that are fitted to a line, giving the relative position and rotation of the two images. The patches are then joined by using a raised cosine blender. This technique is very successful and has also been applied to infrared and X-ray images.


4) *Enhancements to VASARI Since 1993:* We have continued to improve the VASARI scanner since its original completion. We have upgraded the camera A/D to 12 bits/pixel in order to reduce quantization effects and made significant improvements to the color calibration software.

The original VASARI software did a simple LMS fit from the Macbeth chart measurements to find a  $7 \times 3$  matrix, which mapped the seven scanner bands to CIE  $XYZ$ . Unfortunately, a solution that is minimal with respect to errors in  $XYZ$  color space is not minimal visually. We have replaced this with a singular value decomposition to get an initial solution, followed by an iterative scheme using Powell's multidimensional minimization method [18]. This minimizes a function that is a combination of  $\Delta E_{ab}^*$  and the standard deviation of errors. We have also added a preprocessing stage that examines the Macbeth greyscale and adjusts the scale and offset for each band to try to get the black and white points in exactly the right place. This does not use the black or white chart patches as they are the most prone to errors. The results are better overall and can be seen plotted in Fig. 5, where the real colors of the chart can be seen to be uniformly close to the measured values.

One factor that we did not consider carefully enough in the original design is the issue of gloss versus matte finish. The paintings in the National Gallery are almost all finished



(a)



(b)

**Fig. 5.** Plots of (a)  $L^*$  versus  $a^*$  and (b)  $b^*$  versus  $a^*$ , for the real (diamonds) and measured (circles) MacBeth chart values.

with a modern synthetic medium gloss varnish. Although the Macbeth chart is a good spectral match for paintings, it has a very flat matte surface finish. This difference leads to errors in the black point determination, resulting in images of paintings with the black point too low. We have coated our Macbeth charts with a painting varnish to remove this problem.

A significant problem we did not anticipate is that of optical scatter. Imperfections in the camera optics will scatter some of the light from bright areas of the field to dark areas. Since the camera and lights move over the painting, a dark area of the surface might be next to a white patch in one frame and not in an adjacent frame. This can lead to an obvious join




Fig. 6. VASARI scanner in 2000.

line. Scatter is an important factor limiting the accuracy of imaging techniques for color measurement.

We have been able to reduce this problem by observing that the amount of light added to dark areas of the field of view is approximately proportional to the average of the field of view. In other words,  $\text{camera\_response}(x, y) \approx \text{light\_from\_object}(x, y) + \text{mean}(\text{field}) * \text{scatter\_coeff}$ .

We can calculate  $\text{scatter\_coeff}$  by imaging patches of the same reflectivity against backgrounds of different reflectivity. Applying this correction to all the fields in a scan removes visible join lines.

The scanner was refurbished after a laboratory move. A new easel was constructed, which can hold paintings up to 2 m  $\times$  2 m, sliding a 1-m  $\times$  1-m section in front of the scanner. The number of filters was increased to 12, covering 450–1000 nm in 50-nm steps in order to add near-infrared capability. The lighting system was redesigned for the third time, which will double the flux. There were also improvements in the storage of the images due to the Methodology for Art Reproduction in Colour (MARC) and Accurate Colour High Resolution Recording and Reproduction (ACOHIR) projects. Fig. 6 shows the VASARI scanner as it appeared in 2000.

## II. MARC PROJECT

From 1993 to 1996, the European Community-supported methodology for art reproduction in color (MARC) project [19], [20] applied VASARI results to the printing of high-quality art books. We had produced a system capable of accurate color imaging so now the challenge was to complete the chain to the printed book. Within the project, a new scan-back camera was designed and built by Lenz and Lenz at the Technical University, Munich, Germany [21], while the high-level

camera software was developed in London by the authors. The MARC camera used a similar masked CCD to that in the ProgRes camera, but with a combination of macro and micro positioning to cover the image plane of a 6-cm  $\times$  6-cm camera at resolutions of up to 20 000  $\times$  20 000 pixels.

The MARC camera used a RGB color sensor and a 12-bit A/D converter, making it less accurate than the VASARI scanner because of the reduction in the number filters. Offsetting this disadvantage was its portability, ability to image objects of any size, ability to image three-dimensional (3-D) objects and its similarity to a conventional film camera, making the device familiar to photographers. It was also much faster than VASARI, taking around 20 min for a 10 000  $\times$  10 000 scan.

The MARC camera software used color calibration techniques similar to the VASARI scanner. A Macbeth chart was placed in front of the object prior to each session and readings from the chart used to calibrate the images to CIE Lab color space. With HMI (metal halide) arc lamps, the average color error on the Macbeth chart was around  $2 \Delta E_{ab}^*$  units. Fig. 7 shows the MARC camera. One stepper motor can be seen projecting from the top. The conventional camera body provided an easy way to roughly frame the painting using the popup viewfinder.

Crosfield Electronics, one of the MARC partners, developed a package for calibrating and printing to a conventional four-color offset press. They also carried out experiments on seven color printing to see if it provides a wider gamut. These helped us to set the CIE Lab quantization, which is still in use in VIPS as described below. To demonstrate the technology, the MARC project produced a colorimetric catalogue of Flemish Baroque painting [22]. It was a difficult task to




Fig. 7. MARC camera.

calibrate the press in Switzerland and prevent the operators from adjusting the colors, but we managed to produce the book with no conventional proofing.

### III. VIPS IMAGE-PROCESSING LIBRARY

We designed and implemented our own image-processing system in the course of the VASARI and MARC projects. We needed an image-processing package capable of processing very large (around 1 GB) colorimetric images on rather modest hardware. The first Sun workstation we used for development only had a 16-MHz 68 020 processor and 20 MB of random access memory (RAM), yet could still handle large images. The Unix memory map system call is used to give the image-processing functions a virtual-memory mapped image, leaving Unix to handle the paging in of data from the file. This means the VIPS image file format was directly accessed from C in an efficient way with RAM being used by Unix to cache used memory pages automatically. The simple VIPS file format uses a 64-B header and usually uncompressed images with any C type per pixel.

VIPS [23], [24] has been used in all subsequent projects. It has functions to handle color spaces as well as the usual image algorithms, with automatic parallel processing on symmetric multiprocessor machines. The applications described below were implemented using ip, the VIPS graphical user interface (GUI) developed at the National Gallery. VIPS is now an open source project.<sup>1</sup>

The quantized CIE Lab color format (called LABQ in VIPS) stores 10 bits for  $L$ , 11 bits for  $a^*$ , and 11 bits for  $b^*$  so that one sample can be stored in 32 bits and a fine quantization preserved. Conversion to an 8-bit CIE Lab TIFF file is also possible for compatibility with printers and other packages. JPEG compression of the CIE Lab data was also tested as this exploits properties of vision by reducing the spatial resolution of the color channels by 50%.

### IV. APPLICATIONS OF VASARI AND MARC

We have split our imaging work into two parts. The VASARI scanner is used for technical studies of paintings and the MARC camera is used by the photographic department at the National Gallery as a replacement for film photography.

<sup>1</sup>Available at [www.vips.ecs.soton.ac.uk](http://www.vips.ecs.soton.ac.uk).

The VASARI scanner has now imaged around 100 paintings. We are just finishing work on an improved version of the MARC camera: the new camera has a  $1300 \times 1024$  CCD array, Peltier cooling, piezo micropositioning, and improved mechanical positioning. The MARC II can scan, assemble, and calibrate a  $10\,000 \times 10\,000$  image in about 1 min. By June 2001, the National Gallery had imaged the 1000 most important paintings in the collection with the MARC II camera. The camera is available from Lenz and Lenz in Munich, Germany, and is also used in the U.S. Library of Congress, Washington, DC.<sup>2</sup>

1) *Detecting Color Change in Paintings:* In pursuit of our original research aims, we have been using the VASARI scanner to make images of paintings as part of our long-term studies of color change. In the mid 1990s, before sufficient time had elapsed for us to compare images made with the scanner, we compared the color data from the VASARI system with the colorimetric data from the Wright spectrophotometer recorded during the 1970s and 1980s [25]. We developed a procedure whereby the photographic negative that had been made by the Wright spectrophotometer was used to locate the exact area of the digital image corresponding to the reading from the spectrophotometer. The 35-mm negative was digitized along with a 35-mm negative made at the same time that indicated the shape and intensity of the light beam projected onto the surface of the painting by the spectrophotometer optics. These were used to create a mask that defined exactly each of the portions of the painting that had been measured. We then located the corresponding region in the CIE Lab image of the painting made with the VASARI scanner and calculated the average  $L^*$ ,  $a^*$ , and  $b^*$  values and, thus, the  $\Delta E_{ab}^*$  for each measurement point between mid 1980s and mid 1990s. The high spatial resolution of the images from the scanner meant that the area covered by each spectrophotometer reading contained many hundred pixels, providing reliable average values. Generally, we were confident that where the difference between measurements was less than  $4 \Delta E_{ab}^*$  units, no color change had occurred. We identified a number of potential errors, not least, light scatter and glare when measuring the dark colors using the imaging technique. Finally, we noted a few areas of potential color change, that would be reexamined later to determine if this was a real change or the result of a “one off” measurement error.

Now that some years have passed since the first VASARI color images were made, we have been able to make comparisons between colorimetric image of the same paintings recorded at different times. A group of paintings imaged in 1999 have been compared with images made during the period of 1993 to 1996. We have recently described the detailed procedure for this analysis and discussed the inherent errors in [26]. A critical factor in accurate assessment of change was the introduction of a procedure whereby differences in the geometry of the images made at different times could be corrected to allow exact superposition of the “before” and

<sup>2</sup>An interesting cycle of events since the Washington scanner for the U.S. Constitution was an inspiration in VASARI.



“after” images. These differences, which arise from changes in the imaging equipment and the imaging setup, cause differences in scale, rotation, and perspective. A transform of the type

$$x' = x + c_1 + c_2x + c_3y + c_4xy$$

was used to adjust one image with respect to the other, calculating the eight constants ( $c_1$  to  $c_4$  for both  $x$  and  $y$  directions) so that the average differences between the two images was minimized over the whole of their surfaces by selecting many tie points on the images. Fortunately, as the images being compared were sufficiently similar, it was possible to use entirely automatic procedures to derive the parameters to a high degree of accuracy [27].

We have discovered no significant changes in color that can be attributed to color alterations in the seven paintings examined so far, which is encouraging from a conservation point of view. We have, however, discovered that the techniques is sufficiently sensitive that it picks up changes caused by dust settling on the surface of the painting (or being wiped off). Other differences in the images “before” and “after” several years on display can be attributed to slight differences in lighting angle that might cause some specular reflections, shadows from integral frames or impasto, optical flare, or very small misalignments along high-contrast edges caused by the limitations of the resampling of the images.

Although no significant color changes have been detected in the paintings examined to date, a program of continuing measurement is in place.

2) *Detecting Damage to Paintings During Transportation:* Another core objective in building the VASARI scanners was to allow images to be made before and after paintings were transported to exhibitions. Our colleagues at the Doerner Institute have conducted the research in this area and a general description of the methodology used in these studies has been given elsewhere [28]. In a number of studies, it has been shown that small changes can be detected after paintings have travelled to loan exhibitions [21]. More recently, the improved technique for image registration described above has also been applied to images made before and after transportation [29]. This combined with faster image acquisition using the MARC camera has increased the number of studies it has been possible to make. The changes detected ranged from loss or gain of dust particles on the painting surface (as found in the color change study) to displacement of paint particles. Most of the displacements occur at or near the edge of the painting; canvas fiber deformations also occur near the edge of the painting.

As in our color change study, alterations in shadow and gloss create artifacts that produce differences that are clearly not attributable to physical changes. A more severe problem was that some 3-D changes, such as blisters and loosening of tiny particles were hardly detectable. The main advantage of the procedure, however, lies in the computer-aided comparison of two images. In this application, the digital techniques are clearly superior to a simple comparison by eye.

3) *Conservation Documentation:* Colorimetric images of paintings made at each stage of a complicated conser-

vation treatment have been used to provide an accurate permanent record of the change in the appearance of a painting. Although the main purpose of such documentation is to give accurate colorimetric information about the painting at a particular point in its history, the images that have been created during conservation treatments have found other uses. The quality of the images is sufficient that, combined with the recent availability of large format printers at the National Gallery, we have been able to provide restorers with life size prints of the painting. For example, during the recent cleaning of *The Dead Christ Supported by Two Angels* by Crivelli (National Gallery, No. 602), we made images at seven stages during the restoration. Part of the painting, which had been damaged at an unknown point in its history, had been skilfully repainted in the nineteenth century, but, unfortunately, this repainted area had discolored with time. The restorer removed the old repaint, but wished to reconstruct the missing area along the lines of the nineteenth century depiction. With a full-size print of the image alongside her easel, she was able to base the new restoration on the old outline.

We have also made some preliminary studies of using colorimetric images to study the effect on the color of paintings caused by the removal of dirty varnish layers and the possibility of predicting the appearance of a painting with a discolored varnish after it has been cleaned [30]. To do this, we made images of several paintings before and after cleaning. We divided the CIELab color space into cubes of five units ( $L^*$ ,  $a^*$ , or  $b^*$ ) and allocated each pixel to the appropriate cube (color group) within our quantized color space. We then determined, from the “after” image, the color group corresponding to the cleaned color for each pixel. This gave a set of vectors that defined the effect of cleaning on the different colors in the painting. Once we had data for a number of paintings, we attempted to use the model developed on one painting to clean (artificially) one of the other paintings or an as yet uncleaned painting. This met with mixed success, as the degree of yellowing and dirt impregnation varied from varnish to varnish. We investigated the use of small cleaning tests on paintings to determine which of the cleaning models was most appropriate and concluded that this provided the best way to proceed if colorimetric data were to be used. We are currently investigating the use of spectral information to simulate artificial cleaning, combined with imaging ellipsometry to determine the varnish depth across the surface of the painting.

4) *Superposition of Images:* Many museums routinely use infrared reflectography to examine the preparatory drawings beneath the surface of a painting and X-radiography to investigate the physical structure of a painting and changes made during its execution. One of the problems facing art historians and conservators is how to marry up features in the visible image with those beneath the surface. The first hurdle to be overcome is that most museums use lead sulphide-based tube cameras (vidicons) to produce infrared images. The resulting images often suffer from severe barrel or pincushion distortion and the average grey level in a succession of images can be very different as

the detector heats up during the course of session. We have applied software developed in the VASARI project to the computer assembly of infrared reflectograms. First, working with our colleagues in Munich, we have applied the transformations described earlier to reduce the geometric distortion in the images. In this case, an image of a regular grid made with the infrared camera is compared by the software to a perfect grid with the same frequency. The resulting transform is used to correct all the images prior to assembly. To reduce the problem of noise and temperature drift, we have attempted to stabilize the temperature of the tube, either by Peltier cooling or, more simply, by increasing the air flow across the surface of detector [26]. After the reflectogram subimages have been made into a mosaic, the software also performs a global LMS solution using pair overlaps to minimize any differences in grey level in the final image. These reflectograms are easier to make and are of a higher quality than was possible using photographic techniques [17].

Another useful application has involved using a visible-region image to tint an infrared reflectogram. This provides a clearer picture of the relationship between the underdrawing (detected by the infrared camera) and the final painting, making comparisons between draft and final states more straightforward [31]. This is because the visual system uses the color areas as cues to recognize the objects as a whole; the “added” fingers on a hand, for example, do not naturally form part of the rest of the arm.

We have also been developing methods of making X-ray mosaics from the individual large-format photographic plates that are currently produced at the National Gallery. This project has been slower getting off the ground than the infrared imaging project because we have only recently obtained a scanner with sufficient bit depth to capture the information from X-rays and large enough to accommodate the 12-in  $\times$  14-in originals. In addition, we have had to refine the mosaic assembly software used for infrared images to handle the X-ray images. While the infrared images (like those made by the VASARI scanner) generally only require one image to be translated with respect to the other in order to make a join, the X-radiography process is less stringently controlled, so that the images may be rotated or skewed with respect to each other. Again, once we have the X-ray images in the same file format as our visible and infrared images, we have been able to make comparisons of surface and internal features by combining information from the different spectral bands in a single visible image, a topic that will be the subject of a future publication.

Finally, we have been using the resampling techniques mentioned in the context of infrared image correction to allow us to make comparisons between different versions of paintings that might have derived from the same common drawing or cartoon. This is a common task for the restorer or art historian that, up to now, was achieved by making life-size tracings on large acetate sheets and overlaying these sheets to determine differences or similarities in the scale and position of the composition. Given suitable digital images and accurate dimensions for both paintings, this

type of manipulation can be performed with the computer, removing the need to make tracings, which can be inexact and poses a small, but finite risk to the painting. For example, a painting of *Salome* by Giampietrino (National Gallery, No. 3930) shows Salome holding the head of John the Baptist on a platter. Superposition of this figure onto that of Cleopatra in another painting by Giampietrino, now in the Louvre, showed that the same drawing must have been used for both paintings [32]. This type of comparison has now become sufficiently routine that we have developed a small workspace within our image-processing package specifically to handle this operation.

5) *Visualization of Color Change*: The appearance of paintings changes with time. Many of these changes, for example, darkening of the binding medium, occur across the entire surface of the painting. There are other color changes, however, that affect only part of the image and can lead to a loss of meaning in the painting as the balance between colors is irrevocably altered. Because, ethically, no reconstruction of color can be attempted during the treatment of a painting itself, we have been investigating the use of image-processing techniques on high-quality digital images of paintings in which dramatic color changes have occurred to produce an impression of how they may originally have appeared.

Two paintings in which fairly straightforward color changes have occurred are Strozzi’s *Annunciation* (National Gallery, No. 1406) and Veronese’s *Vision of St. Helena* (National Gallery, No.1041). The sky in the Veronese was painted using the blue cobalt-containing pigment smalt. Unfortunately, with time, this pigment produces a brown hue as the blue particles lose their color and the oil binding medium is turned brown by reactions catalyzed by the cobalt. Using samples of smalt prepared according to a traditional recipe, we were able to gain information on the color of the original sky, although we found that all the smalt we produced was more intense in color than the original materials, as we could not easily take into account the poor quality of materials used in the sixteenth century. We first created a mask covering the area of the sky (by selecting the region according to its  $a^*$  and  $b^*$  values and then editing the mask edges by hand). As the original lightness of the sky seemed to have been well preserved despite the shift in hue, we made no attempt to alter the  $L^*$  values. Our experiments with original materials indicated that most smalt samples had a very similar hue and that the saturation (chroma) was closely related to lightness. Accordingly, we substituted a fixed hue for all the pixels covered by the mask and scaled the chroma so that it was greatest in the darkest areas (setting the maximum and minimum chroma values by measuring areas of smalt that have survived in other paintings). The end result is a sky that now appears blue, but that has not lost the subtle changes in lightness and chroma from the horizon to the top of the painting and in which the wispy clouds painted by Veronese can clearly be seen (see Figs. 8 and 9).

In the *Annunciation* the garments of two principal figures, the Virgin Mary and the Angel Gabriel, were painted using fugitive red lake pigments mixed with varying amount of lead



Fig. 8. Veronese before processing.

white. The main problem encountered here was that the loss of color from the mid- and light-tones is so great that it is no longer to distinguish the regions of original highlight from the mid-tones. For the reconstruction, we relied on measurements made on a series of mixtures of the original red pigment mixed with differing amounts of lead white, before and after it had been subjected to accelerated light ageing [33].

Recently, we have been studying the changes in a number of different pigments in a single painting, Campaña's *The Conversion of the Magdalen* (National Gallery, No.1241). The painting depicts many figures, originally in brightly colored draperies, many of which have deteriorated to different shades of brown. As the original color of these areas could only be deduced by scientific analysis of the pigments, this painting proved a particularly suitable candidate for color reconstruction on a digital image. A full report on these studies is in preparation.

6) *Pigment Identification*: We have investigated using VASARI multispectral images to identify pigments in paintings [34]. Although we were able to successfully identify some pigments, the spectral resolution and range of VASARI at the time we did the study was not sufficient for practical conservation use.

The seven-band images give enough information for a good CIE Lab measurement, but there is not enough detail



Fig. 9. Veronese after processing.

to allow spectra to be identified accurately. Pigment spectra are usually rather similar, often very small intermediate maxima or “shoulders” are the only features distinguishing two materials with the same color but different chemistry. Identification is also hampered by the frequent use of complicated pigment mixtures.

7) *Printing*: MARC gave us a better understanding of printing from the colorimetric images. We now use calibrated large-format inkjet printers donated by Hewlett-Packard (2500CP) to produce color-accurate life-size prints of paintings. The prints have been used as reference images during conservation treatment and to print facsimiles of paintings for loan to other galleries when the original painting is too fragile to be transported. Their color accuracy (MacBeth average around  $3 \Delta E_{ab}^*$ ) and the ability to print details on demand has been of use inside the National Gallery as well as to the public.

## V. VISEUM—PUTTING THE IMAGES ON THE WEB

In 1996, the Viseum project worked on a client-server system [35] to make the large images produced in previous work accessible over the Internet. Our solution used the internet imaging protocol (IIP) with extensions to allow a Java client in a normal web browser to fetch sRGB image tiles ( $64 \times 64$  pixels) on demand in JPEG format for display.



Fig. 10. High-resolution image shown in web viewer.

IIP standardized the way that programs request images from a server so that an IIP-aware browser can look at servers from different vendors. Thus, it was possible to browse any size image over the web.

Although Photo-CD was capable of storing a multiresolution pyramid of images, it was proprietary and inflexible (also true of the later Flashpix). Rather than evolving the VASARI format used in VIPS, the standard TIFF format was used. A multipage TIFF file was used with full, half, quarter, etc. resolutions stored as tiled JPEG images. This means that the file is 30% bigger, but makes access much faster. By fetching only the tiles necessary for the specific view required at any one time the system became usable over the Internet at low bandwidths. An optional display color profile (International Colour Consortium) server maintains profiles on an internet protocol (IP) address basis so that images can be transformed specifically for the user's monitor rather than to sRGB. In general use, however, sRGB versions are stored for simplicity, as this standard became more common. It is major differences in monitor gammas that actually appear to cause the most problems as the eye sees differing amounts of contrast in the light or dark areas depending on the monitor. SRGB sets a monitor gamma of 2.2. Fig. 10 shows the Viseum-viewer browsing the low-resolution view of Renoir's *The Umbrellas* (National Gallery, No.3268) with Fig. 11 showing the full resolution detail of this 7275 × 9982 pixel image. Uncompressed it




Fig. 11. Closeup of the full 7295 × 9982 pixel image.

occupies 290 MB and the JPEG-TIFF is 50 MB. This low level of JPEG compression does introduce visible artifacts.

## VI. ACOHIR

The European project ACOHIR in 1996 applied the color calibration from MARC with the Viseum image viewer to produce higher quality object views of objects [36]. These were captured with a high resolution Kodak digital camera and a motorised turntable. Objects could be seen from a wide range of angles at higher resolutions than available previously using Apple's QuickTime-VR, for example. The systems were used in archaeological and commercial applications and a scanner was made for the Louvre using a firewire-linked DCS560 camera capturing 3 × 36 views of 3000 × 2000 pixels each. The average  $\Delta E_{ab}^*$  over the MacBeth chart was still around three, which was good for a commercial RGB system. As in Viseum, the images were transmitted to the client as sRGB. In order to maintain consistency with different monitors, users were expected to set up their monitors to be as close to sRGB as they could.

## VII. CONCLUSION

In ten years, we have implemented and continued to develop different imaging systems linked by a common image file format and image-processing library. Digital imaging has become an established part of the National Gallery's activities and is beginning to replace conventional photography. Indeed, over the period of 1999–2002, the

National Gallery has used two MARC cameras (an improved version of those developed in the synonymous project) to digitize all its major paintings. If the second phase of this project goes ahead, high-resolution (10 000 pixels along the longest dimension of the painting) colorimetric images of all the paintings in the collection will be available on the Gallery's server. These images will be the basis of all electronic and traditional publications by the Gallery, from information sheets to scholarly catalogues. The VASARI system's main use continues to be long-term monitoring of the condition of paintings and the technical analysis/documentation of the National Gallery's collection. We are still actively developing the system, looking for improvements in image quality, spectral range, and spectral reconstruction techniques.

Although the origins of these projects were in the scientific analysis of paintings, increased awareness of the usefulness of electronic images has led to a range of applications in addition to those we originally envisaged. Advances in CCD resolution have still not outdated the mosaicing techniques we developed. Color calibration is still an active area of research as it is almost impossible to achieve an exact solution. We expect to increase precision further in the future using the latest advances in color science.

#### ACKNOWLEDGMENT

The authors would like to thank Barco for advice and monitor models for color display as well as sponsorship for their Calibrator. The MARC and VASARI projects were funded by the European Commission's ESPRIT program. VASARI involved: the Doerner Institute, Munich, Germany; Brameur Ltd., U.K.; Birkbeck College, London, U.K.; Telecom Paris/ENST; Thomson-CSF LER, Rennes, France; and TÜV-Bayern, Munich, Germany. The ESPRIT III Project MARC involved: Thomson Broadband Systems; the National Gallery; Birkbeck College; the Doerner Institute; Crosfield Ltd., U.K.; Hirmer Verlag, Germany; and Schwitler, Switzerland. The new VASARI scanner GUI was written by R. Pillay while at the National Gallery funded by Hewlett-Packard Laboratories.

#### REFERENCES

- [1] W. D. Wright, "A mobile spectrophotometer for art conservation," *Color Res. Applicat.*, vol. 6, pp. 70–74, 1981.
- [2] M. P. Wassall and W. D. Wright, "A special purpose spectrophotometer," in *AIC Congress Proceedings*. London, U.K.: Adam Hilger, 1973, pp. 469–471.
- [3] L. Bullock, "Reflectance spectrophotometry for measurement of color change," *Nat. Gallery Tech. Bulletin*, vol. 2, pp. 49–55, 1978.
- [4] G. Thomson and S. Staniforth, "Identification and measurement of change in appearance by image processing," in *Science and Technology the Service of Conservation*, N. S. Brommelle and G. Thomson, Eds. London, U.K.: Int. Inst. for Conservation, 1982, pp. 159–161.
- [5] D. Saunders, "Color change measurement by digital image processing," *Nat. Gallery Tech. Bulletin*, vol. 12, pp. 66–77, 1988.
- [6] *Recommendations on uniform color spaces, color difference equations, psychometric color terms*, Supplement No. 2 to CIE Publication No. 15 (E-2.3.1), 1978.
- [7] K. Martinez, J. Cupitt, and D. Saunders, "High resolution colorimetric imaging of paintings," in *Proc. Soc. Photo-Optical Instrum. Eng.*, vol. 1901, Feb. 1993, pp. 25–36.
- [8] D. Saunders and J. Cupitt, "Image processing at the national gallery: The VASARI project," *Nat. Gallery Tech. Bulletin*, vol. 14, pp. 72–85, 1993.
- [9] D. Saunders and A. Hamber, "From pigments to pixels: Measurement and display of the color gamut of paintings," in *Proc. Soc. Photo-Optical Instrum. Eng.*, vol. 1250, Feb. 1990, pp. 90–102.
- [10] A. Burmester, J. Cupitt, H. Derrien, N. Dessipris, A. Hamber, K. Martinez, M. Müller, and D. Saunders, "The examination of paintings by digital image analysis," in *Proc. 3rd Int. Conf. Non-Destructive Testing, Microanalytical Methods and Environmental Evaluation for Study and Conservation of Works of Art*, Viterbo, Italy, 1992, pp. 201–214.
- [11] A. R. Calmes and E. A. Miller, "Registration and comparison of images obtained at different times for ageing studies of the U.S. constitution," in *Proc. Soc. Photo-Optical Instrum. Eng.*, vol. 901, 1988, p. 61.
- [12] G. Deconinck, "Multi-channel images: Acquisition and coding problems," M.Sc. thesis, Katholieke Universiteit Leuven and Ecole Nationale Supérieure de Télécommunications, Leuven, Belgium, 1990.
- [13] P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, "Image capture: Simulation of sensor responses from hyperspectral images," *IEEE Trans. Image Processing*, vol. 10, pp. 307–316, Feb. 2001.
- [14] R. Lenz, "Calibration of a color CCD camera with 3000 × 2300 picture elements," in *Proc. Soc. Photo-Optical Instrum. Eng.*, vol. 1393, 1990, pp. 104–111.
- [15] C. S. McCamy, H. Marcus, and J. G. Davidson, "A color-rendition chart," *J. Appl. Photo. Eng.*, vol. 2, no. 3, pp. 95–99, 1976.
- [16] A. Hård and L. Sivik, "NCS—natural color system: A Swedish standard for color notation," *Color Res. Applicat.*, vol. 6, pp. 129–138, 1981.
- [17] A. R. Billinge, J. Cupitt, N. Dessipris, and D. Saunders, "A note on an improved procedure for the rapid assembly of infrared reflectogram mosaics," *Studies Conservation*, vol. 38, pp. 92–98, 1993.
- [18] R. P. Brent, *Algorithms for Minimization Without Derivatives*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [19] J. Cupitt, K. Martinez, and D. Saunders, "Methodology for art reproduction in color: The MARC project," *Comput. History Art*, vol. 6, pp. 1–20, 1996.
- [20] H. Derrien, "MARC, a new methodology for art reproduction in color," *Inform. Services Use*, vol. 13, no. 4, pp. 357–369, 1993.
- [21] R. Lenz, R. Beutelhauser, and U. Lenz, "A microscan/macroscan 3 × 12 bit digital color CCD camera with programmable resolution up to 20 992 × 20 480 picture elements," in *Proc. Commission V Symp. Close Range Techniques and Machine Vision*, vol. 30, Melbourne, Australia, 1994.
- [22] A. Burmester, L. Raffelt, K. Renger, G. Robinson, and S. Wagini, *Flämische Barockmalerei: Meisterwerke der Alten Pinakothek München Flemish Baroque Painting: Masterpieces of the Alte Pinakothek München*. München, Germany: Hirmer Verlag, 1996.
- [23] J. Cupitt and K. Martinez, "VIPS: An image processing system for large images," in *Proc. IS&T/SPIE Symp. Electronic Imaging: Science and Technology*, vol. 2663, Very High Resolution and Quality Imaging, 1996, pp. 19–28.
- [24] —, "Image processing for museums," in *Interacting With Images*, L. MacDonald and J. Vince, Eds. New York: Wiley, 1994, pp. 133–147.
- [25] D. Saunders, H. Chahine, and J. Cupitt, "Long-term color change measurement: Some results after twenty years," *Nat. Gallery Tech. Bulletin*, vol. 17, pp. 81–90, 1996.
- [26] D. Saunders, A. Burmester, J. Cupitt, and L. Raffelt, "Recent applications of digital imaging in painting conservation: Transportation, color change and infrared reflectographic studies," in *Tradition and Innovation: Advances in Conservation*, A. Roy and P. Smith, Eds. London, U.K.: Int. Inst. for Conservation, 2000, pp. 170–176.
- [27] R. Lenz, "Ein Verfahren zur Schätzung der Parameter Geometrischer Bildtransformationen," (in German), *Nachrichtentechnische Berichte Bd. 15*, Tech. Univ., Munich, Germany, 1985.
- [28] A. Burmester and W. Wei, "All good paintings crack: Nondestructive analysis of transport damage of paintings using digital image processing," in *Proc. 4th Int. Congr. Non-Destructive Testing of Works of Art*, Berlin, Germany, 1994, pp. 114–126.
- [29] A. Burmester and M. Müller, "The registration of transportation damages using digital image processing," *Z. Kunsttechnol. Konservierung*, vol. 6, pp. 335–345, 1992.

- [30] H. Chahine, J. Cupitt, D. Saunders, and K. Martinez, "Investigation and modeling of color change in paintings during conservation treatment," in *Imaging the Past: Electronic Imaging and Computer Graphics in Museums and Archaeology*, T. Higgins, P. Main, and J. Lang, Eds. London, U.K.: British Museum, 1996, Occasional Paper 114, pp. 23–34.
- [31] D. Saunders and J. Cupitt, "Elucidating reflectograms by superimposing infra-red and color images," *Nat. Gallery Tech. Bulletin*, vol. 16, pp. 61–65, 1995.
- [32] L. Keith and A. Roy, "Giampietrino, Boltraffio and the influence of Leonardo," *Nat. Gallery Tech. Bulletin*, vol. 17, pp. 4–19, 1996.
- [33] J. Kirby, D. Saunders, and J. Cupitt, "Colorants and color change," in *Early Italian Paintings Techniques and Analysis*, T. Bakkenis, R. Hoppenbrouwers, and H. Dubois, Eds. London, U.K.: Archetype, 1996, pp. 60–66.
- [34] J. Thoma, J. Cupitt, and D. Saunders, "An investigation of the potential use of visible-region multispectral imaging for pigment identification in paintings," in *Proc. Color Image Science Conf.*, Apr. 2000, pp. 95–106.
- [35] K. Martinez, J. Cupitt, and S. Perry, "High resolution colorimetric image browsing on the web," in *Comput. Netw.*, vol. 30, 1998, pp. 399–405.
- [36] K. Martinez, S. Perry, and J. Cupitt, "Object browsing using the internet imaging protocol," *Comput. Netw.*, vol. 33, pp. 803–810, 2000.



**John Cupitt** received the B.Sc. degree in computer science and the Ph.D. degree in theoretical computer science from the University of Kent, Kent, U.K., in 1986 and 1989, respectively.

He is currently with the Scientific Department of the National Gallery, London, U.K., where his research includes conservation applications of digital imaging.



**David Saunders** has received the B.Sc. degree in chemistry and the D.Phil. degree in organometallic chemistry from the University of York, Heslington, U.K., in 1981 and 1984, respectively.

Since 1985, he has been with the Scientific Department of the National Gallery, London, U.K., investigating the deterioration of artists' materials and the use of digital imaging in the technical examination of easel paintings.



**Kirk Martinez** (Member, IEEE) received the B.Sc. in physics from the University of Reading, Reading, U.K., in 1983 and the Ph.D. degree in image processing from University of Essex, Wivenhoe Park, U.K., in 1990.

He was an Arts Computing Lecturer with the University of London developing applications of image processing and imaging to art. He is currently with the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K. His current research interests

are content-based retrieval of art and museum application of augmented reality.



**Ruven Pillay** received the M.Sc. degree in computer science from the University of Edinburgh, Edinburgh, U.K., in 1994 and the M.Phys. in physics from the University of Manchester, Manchester, U.K., in 1996.

He was a Researcher with the National Gallery, London, U.K., and is currently with the Central Research Laboratories of the Louvre, Paris, France.