15 The Art of Virtual Restoration: Creating the Digital Image Archive of Medieval Music (DIAMM)

Andrew Wathey[1], Margaret Bent[2], and Julia Craig-McFeely[1]

[1]Royal Holloway Bedford New College
University of London
London, UK
a.wathey@lara.rhbnc.ac.uk
julia.craig-mcfeely@music.ox.ac.uk

[2]All Souls' College
University of Oxford
Oxford, UK
margaret.bent@all-souls.ox.ac.uk

Abstract

The Digital Image Archive of Medieval Music is a new resource for scholars which aims to digitize, archive, and make available images of British manuscripts of medieval polyphonic music, and to develop techniques of digital image-enhancement, or "virtual restoration," to retrieve lost data or improve the legibility of materials that cannot at present be read. The article surveys the work of the project in digitizing primary sources, in creating the archive and in delivery to the end user. It also examines methods whereby image processing can enhance the examination of these sources. These methods are relevant to any discipline in which the study of original sources is undertaken.

COMPUTING IN MUSICOLOGY 12 (1999-2000), 227–240.

15.1 Introduction

The Digital Image Archive of Medieval Music[1] aims to digitize, archive and enhance images of British manuscripts of pre-Reformation polyphonic music, a rich and neglected repertory that survives mainly in fragmentary and often barely legible sources. The first phase of the project involved the collection of digital images and computer enhancement of fifteenth-century fragments, with a view to publication in facsimile form.[2] In its second phase, the project has expanded to embrace all of the fragmentary and some of the less accessible complete and near-complete manuscripts. Chant is at present excluded, but fragments of non-British manuscripts of polyphony in UK repositories, which are very few in number, will be digitized as opportunity permits. The aims of the project are: to create a new permanent electronic archive of images of manuscripts of British medieval polyphony, to facilitate wider study of their contents and to assure their permanent preservation; to develop techniques of digital image-enhancement (or "virtual restoration") to make legible materials that cannot at present be read; and to develop new and more exact types of manuscript study made possible by digital image technology. Some fifteen manuscripts survive in complete or near-complete form, of which only four date from the fifteenth century and only one (the tenth-century Winchester Troper) is earlier than c.1415-20. All other manuscripts between the Norman Conquest and 1548 are fragmentary; a very few of these are substantial groupings of up to 24 pages, but the majority are single leaves, pairs of leaves or smaller scraps. A total of 354 fragments are known, witnesses to an almost equivalent number of books that once existed complete. Taken together they amount to c.2250 pages, thus adding about 60% to the total number of pages in complete and near-complete manuscripts. Many of these fragments have come to light since 1945: the largest number received their first published notice in the 1950s and since c.1980 discoveries have been made at a rate of about 60-70 leaves per decade. In the last 30 years, the number of manuscript leaves known from pre-Reformation Britain has expanded by over a third, with new finds frequently forcing the re-ordering of a repertory or an individual composer's output.

Most important among the reasons for loss is probably musical obsolesence. Choirbooks were dismembered when their contents fell out of use and their leaves were reused (because parchment was both strong and relatively expensive) to bind books and documents or for other

[1] Work on the first phase of the project began in 1998 with funding from the Humanities Research Board of the British Academy. The second, current phase is the recipient of a Major Research Grant from the Arts and Humanities Research Board, awarded in 1999.

[2] *English Manuscripts of Fifteenth-Century Polyphony in Facsimile, Early English Church Music*, forthcoming.

purposes. Smaller, ephemeral materials, including some part-books but also rolls and loose sheets, were sometimes similarly reused. In consequence, the survival of polyphony is yoked to broader patterns of manuscript survival, to the history of binding, and to shifting attitudes, since the seventeenth century, toward collecting and preservation. Collections unscathed by the mass re-binding programmes of the nineteenth-century, when it was still common to discard the old binding in toto, are usually rich in such fragments; libraries that could afford the cosmetic uniformity of wholesale rebinding are usually less so.[3] Many discoveries have been made—and continue to be made—in smaller collections, in county record offices, cathedral libraries, and some private collections. Total losses are inevitably hard to estimate, but it is clear that even c.350 volumes can represent only a very tiny fraction of what must once have existed.[4]

[3] Most notably the British Museum in London and the University Library in Cambridge; see Wathey 1993.

[4] See, for example, the references to books no longer extant, drawn from medieval inventories and accounts, in Wathey 1988.

Reconstructing the music transmitted by these fragments is hampered not only by physical damage but also by the notational practices of the period. In most polyphony, individual voice-parts were notated separately, in different portions of the page, rather than in "score" with the parts in parallel. Medieval choirbooks were designed for performers to use, and the most common notation placed single voice-parts on facing pages of an opening; single leaves may thus often transmit a whole work, but in only half its voice-parts.

In "score," where all the voices were notated simultaneously, whole compositions rarely fit on one page; in this case a single leaf may transmit all voices but for only part of a work. However, the presence of several shared pieces among the surviving leaves—indicative of a largely common repertory—can mitigate such losses. In more than one case, it has been possible to reconstruct a complete composition from incomplete versions present in two or more fragments.[5]

[5] See, for example, the reconstructions by Bent (1970, 1984).

15.2 Digitizing Primary Sources

The digital medium offers significant advantages over conventional media in storage, reproduction quality and delivery. Digital images are a near-permanent form of record, whereas even carefully stored conventional media deteriorate with time and use. Conventional (analogue) images cannot be copied without loss of data and degradation between original and copy; digital images can be copied exactly (with, as is well known, implications for copyright). Image quality using a high-resolution digital camera significantly exceeds that available from library stills cameras, permitting greater accuracy of focus and recording more detail in both visible and non-visible spectra. Delivery via Web-based systems permits very wide dissemination of these materials, study of which has until now been hampered by their considerable geographic spread.

For the DIAMM project, direct digitization from originals rather than from surrogates (e.g. pre-existing slides and photographs) has been adopted wherever possible, to facilitate the acquisition of the largest images at the highest available resolution. The scanning of surrogates introduces a further step between the original and digital image at which data may be lost, whether by error or the limitations of analogue reproduction. Errors may also arise because the eye is unable to perceive color over the full light spectrum (particularly in the blue range) and has a limited focussing ability. These further levels of definition prove essential in the study and enhancement of damaged sources. More important, the quality of a scan from a surrogate is only as good as that of the target slide or photograph. This is usually well below that obtainable from the original with a high-resolution digital camera.

The 7000 x 7000 pixel array of the PhaseOne Powerphase scanning-back used for DIAMM photography (and similar equipment used for image capture in major libraries) produces images of 49 million pixels at 24-bit color-depth. (A single-shot library stills camera, comparable with conventional professional archive photography, may capture at 2800 x 3500 pixels [9.8 million pixels].)

Iterative focussing based on greatly enlarged "preview" images has produced results of significantly higher definition than conventional (eye-focussed) methods.[6] Capture times vary but are usually about 5 minutes under daylight-balanced, cold fluorescent lights. Images of this size tolerate greater enlargement without "graining" than their analogue counterparts. Moreover, much of the additional information in images of this size is apparent to the naked eye only at high magnification, and its full potential is manifest only when manipulated for restoration or other purposes. Large digital images of this type also promise to facilitate new types of manuscript study, in particular pixel-level analyses of ink and other surface features and of graphic variance in the work of individual scribes.

[6] Some examples and further discussion are available at *www.diamm.ac.uk*.

15.3 Creating the Digital Archive

At the time of writing (mid-2000), high-resolution image capture is available in a relatively small number of major British libraries.[7] Although the number of repositories with this facility is likely to grow with time, image capture for materials elsewhere has been undertaken with equipment purchased specially for the DIAMM project. This mobile digital studio comprises a high-resolution digital camera (a PhaseOne PowerPhase digital scanning-back mounted on a medium format Fuji GX 680 III PRO camera body and lens), copy stand, lighting units (daylight-balanced and UV), laptop computers, and 36Gb external hard drive for storage during *in situ* photography in libraries. Individual images produced range from 80Mb to 144Mb; larger scanning-backs are available but produce images in the 1Gb range, which are as yet impractical to manipulate.

[7] The British Library (London), Cambridge University Library, the Bodleian Library (Oxford), the National Library of Scotland (Edinburgh) and the National Library of Wales (Aberystwyth). Some other repositories have access to 10-megapixel cameras through outside agencies.

Data storage is a major concern in any image archive. Copies on CD-ROM are used for medium-term storage and enhancement and other work within the project team. Two hierarchical file-servers (in different locations), based on linked optical disks and providing redundant backup systems, ensure the long-term safety of the data. This system offers virtually unlimited storage capacity at the highest standards of data preservation, but is suitable only for archiving; short-term retrieval is relatively slow and dependent on the level within the system at which data is stored.

Files are captured, saved, and archived in uncompressed TIFF format to maximize compatibility. Lossless compression is not used, to avoid the risk of minor file corruption. Images incorporate industry-standard color patches and a linear scale to allow accurate calibration in subsequent on-screen or print use. Lossy file-compression formats (principally JPEG) are used only to generate smaller images that are used via the Web or for demonstration purposes.

Intellectual, technical, and administrative metadata are collected in a *Filemaker Pro* database that was developed, using the metadata schema produced for the JISC Image Digitization Initiative (JIDI) as a starting point.[8] Metadata can be exported in a number of formats, including *XML*, Dublin Core, etc. But the database incorporates further manuscript-level data and is usable both in conjunction with the images and as an independent searchable study resource, which will ultimately be

[8] The JIDI feasibility study can be found at *heds.herts.ac.uk/Guidance/JIDI_fs.html.* The JIDI metadata guidelines were themselves derived from the VRA recommendations.

Table 1.

country	UK
county	Worcestershire
city/town	Worcester
archive	Dean and Chapter Archive
source	Add. 68
folio	x 2r
photographer	Julia Craig-McFeely
date created	21/9/1999
capture device	PhaseOne PowerPhase on medium format Fuji GX680III professional body
light source	reflective, fluorescent
lighting	4 x daylight balanced LARN Softscan high-frequency 110W non-dimmable heads
light-source distance	108 cm
file format	TIFF
compression ratio	none
compression type	none
scanner mode	24-bit color
pixel resolution	7000 x 7000
scan resolution	600 dpi
color space	RGB
color management	not applied
film curve	linear response/standard capture conditions
software	*PhaseOne* 3.1.1
scan setting	high-resolution
orientation	landscape
color correction	off
focus	manual
cropping	none
rotation	none
contrast	n/a
brightness	n/a
filters	Tungsten IR
white point	6500°
gamma	1.8
stored size	112.4 MB
filename	GB-WORcAdd68_x_2r.tif
storage site	ftp://..........
image URL	http://..........
thumbnail	[+]

made available online. Extensive data is compiled for each individual image, substantially above that suggested in Dublin Core, to maximize future utility and forward compatibility. A sample is given in Table 1.

15.4 Delivery to the End-User

Web-based dissemination provides end-users with a readily accessible resource, removing the natural inhibition to studying medieval music that has long been imposed by its geographic spread. Publication on CD-ROM, or through Web/CD hybrids, may prove an effective means of disseminating the whole archive when it is complete. Both Web and CD dissemination impose an equipment requirement for the end-user (and accurate color reproduction is crucially dependent on correct monitor calibration). It is likely that conventional publication in hard copy will also be used.

Systems for Web-based dissemination have been established to allow a hierarchy of different levels of access and to provide effective means of copyright protection. Images are normally delivered at 72–96 dpi, sufficient for detailed manuscript study on-screen and allowing download within an acceptable period of time; viewers such as *FlashPix* can be used to provide swift access to versions at higher resolution.

Web-based materials are protected by copyright laws: end-user access requires an individual password and username to access images, granted subject to acceptance of a signed Web-access agreement with specific copyright stipulations. Images are print-disabled and carry a copyright notice. License agreements have been evolved for use with libraries and archives holding materials included in DIAMM, to allow digitization and Web-publication in a copyright-protected environment. The DIAMM portal (at *www.diamm.ac.uk*) offers libraries and archives the opportunity to bring their collections to wider public notice under proper controls.

15.5 Virtual Restoration

The use of digital image-processing techniques to restore semi-legible and illegible material has been viable for some time, but only recently has it become widely available in humanities research, as commercial software increases in sophistication and decreases in cost.[9] As a non-invasive

[9] See, for example, Benton, Gillespie, and Soha (1979).

process, "virtual restoration" avoids the risk of damage that has dogged all physical restoration techniques, many of which (most notoriously the use of galls) damaged the originals to an extent that precludes the use of more recently developed methods. In developing restoration techniques the DIAMM project has deliberately chosen affordable commercially available software (Adobe *Photoshop*), in order that others should be able to replicate or improve on our results. Commercial software also offers support and development investment that many private developers are unable to provide.

Individual pages, even from the same manuscript, frequently present widely differing problems. Writing may be obscured by dirt, glue or overwriting; the written surface may be rubbed, faded or stained. Any leaf from any source may thus require a unique approach, rendering single-algorithm processes inappropriate, and the same is frequently true of discrete portions of the same page. The identification of such areas is essentially an editing decision dependent on specialist paleographical and repertorial knowledge, to the extent that two scholars might potentially identify different sets of marks as dirt to be discarded and ink to be retained.

This issue does not arise with "global" tools and transformations, e.g. color-specific processes applicable to whole pages, where the ink color can be selected in a part of the page where it is clearly identifiable. But (e.g.) in removing staining that is non-uniform or that is very close in color to the underlying ink (in both cases portions of the original text may be lost), the expertise of the restorer in understanding the notation that is obscured is essential. Almost all restoration is a form of editing.

Adobe *Photoshop* was designed mainly for creative graphic work and has become an established tool for both print and the Web. Many of its functions are not relevant to "virtual restoration" but others, developed with different applications in mind, have proved highly effective.

Three tools, with their subroutines, are of particular value in "virtual restoration": level-adjust, the layer function, and the color-range selection tool. These are briefly surveyed below. The restorations cited can be seen at *www.diamm.ac.uk*, with a fuller explanation of current techniques.

Figure 1. Parchment wrapper before and after "virtual restoration" by DIAMM. Private collection of Christopher de Hamel, reproduced by the kind permission of Christopher de Hamel.

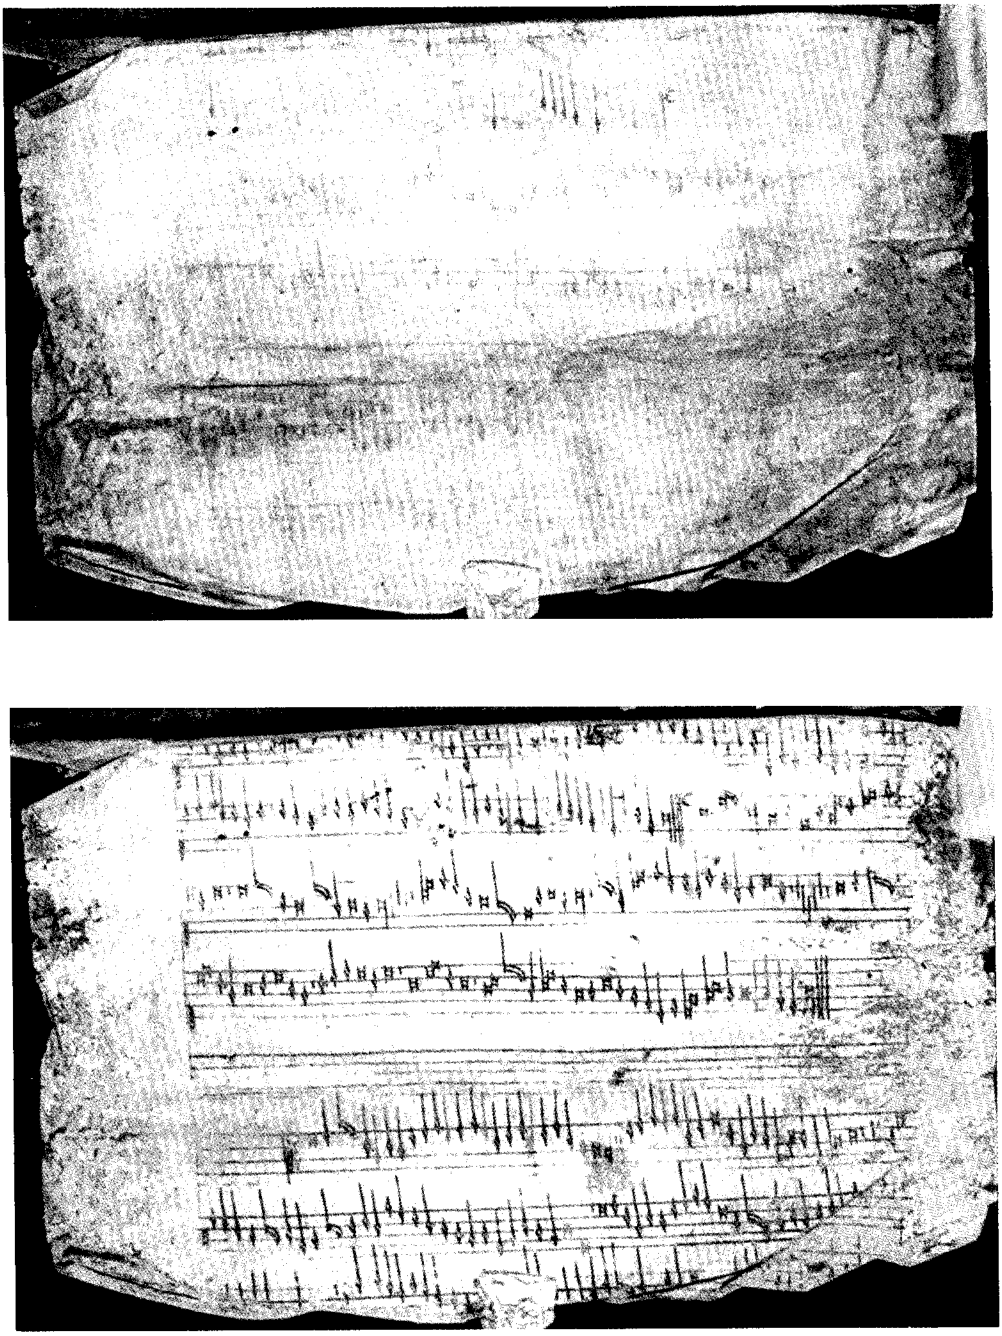

Figure 2. London, British Library, Add. 41340 before and after "virtual restoration" work by DIAMM. ©2000 The British Library. Permission pending.

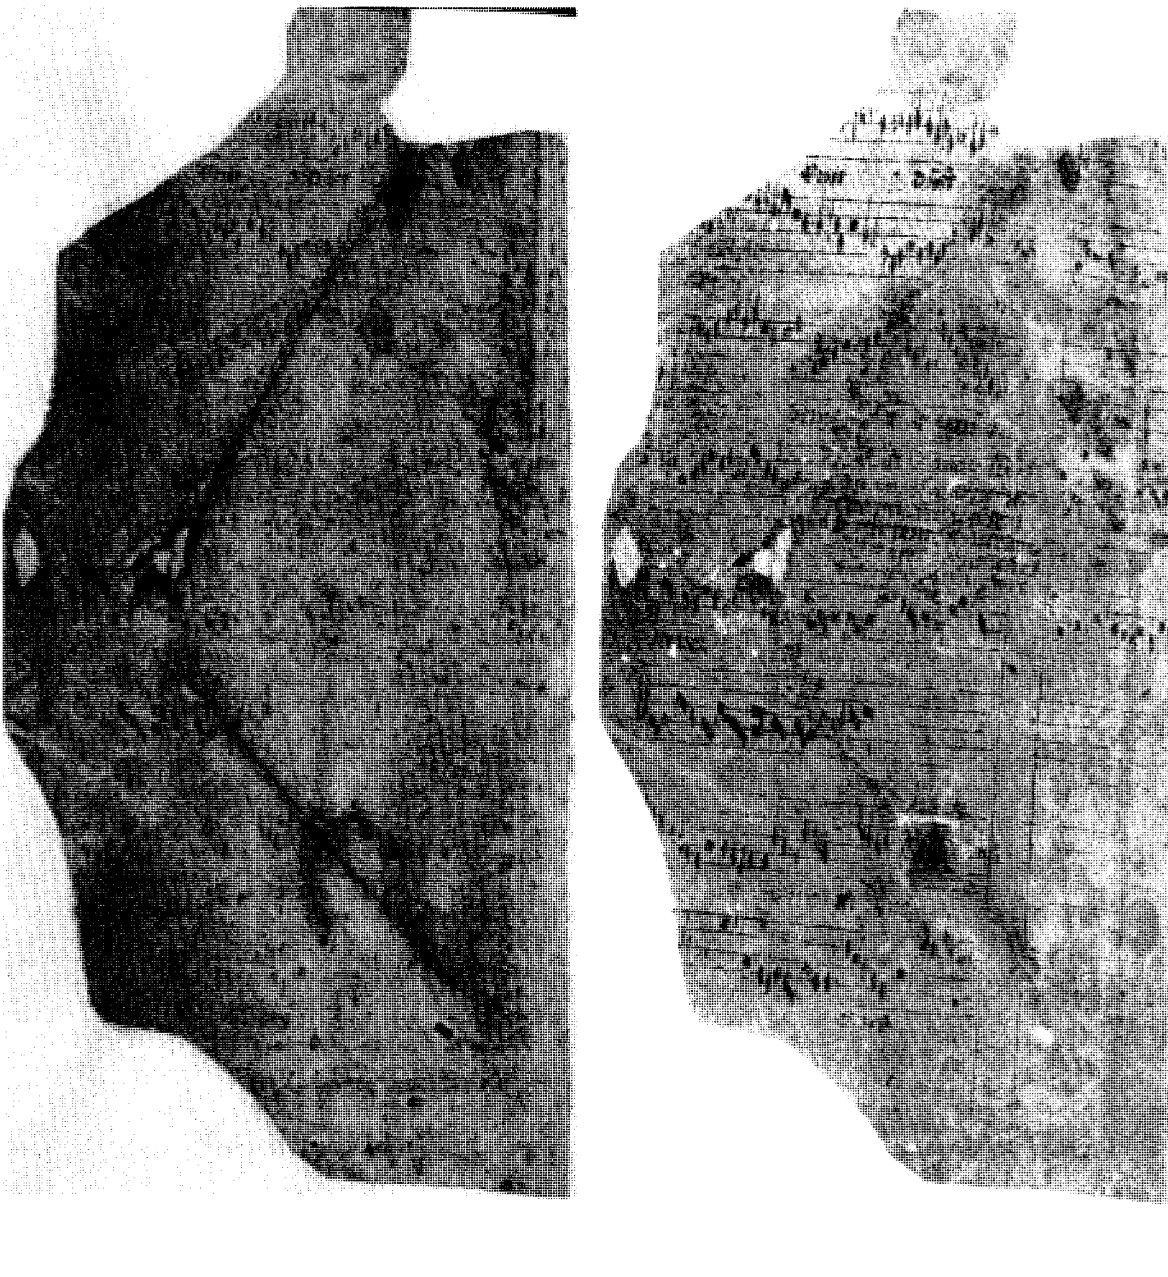

Figure 3. GB-Oxford, Corpus Christi College, MS 144: four stages of restoration by DIAMM on one of the palimpsest folios. ©1999 Corpus Christi College, Oxford, reproduced by the kind permission of the President and Fellows.

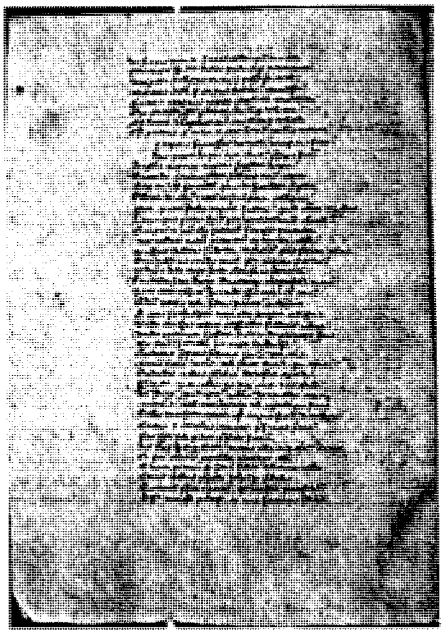

Stage A

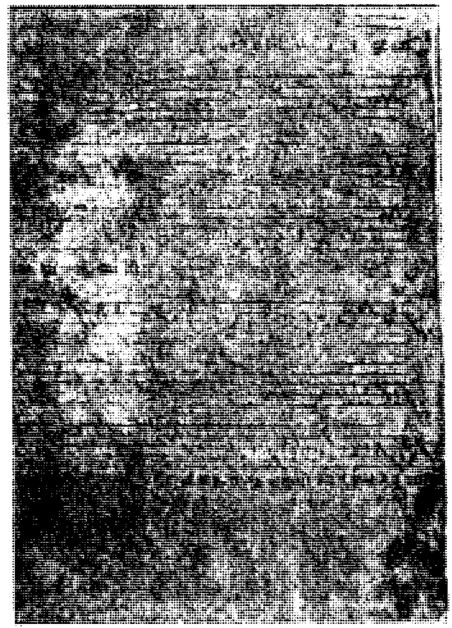

Stage B

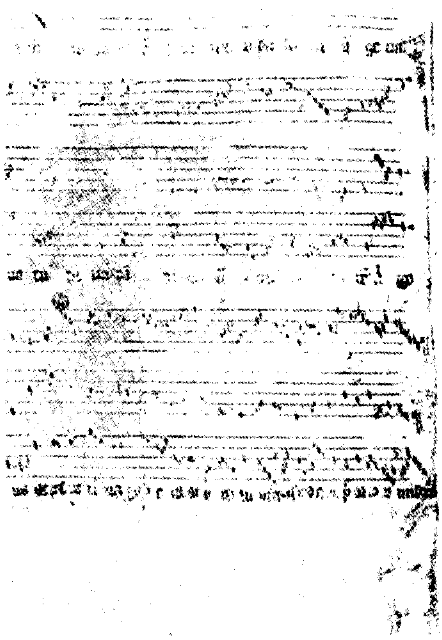

Stage C

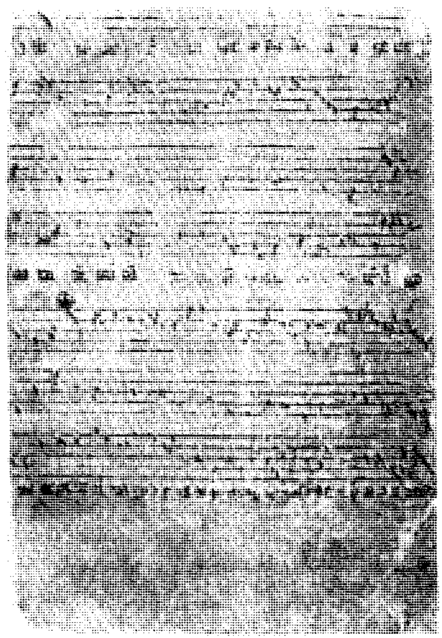

Stage D

1. LEVEL-ADJUST is a basic tool in image-processing software. Originally designed to compensate for over- or under-exposed images, this tool has proved particularly useful in isolation, to improve general legibility, and in combination with other tools. The tool presents a histogram graphing the numbers of pixels in the image from values of 0 (black) to 255 (white). "Sliders" and other tools allow the user to redefine the black, white, and midpoints, or to adjust them up or down. In a relatively contrast-less source, with color response in the middle of the graph only, movement of the sliders inwards can be used to spread a narrow color range, darkening the deeper and lightening the paler colors. New colors are not generated; each color-value is separated from its nearest value by a greater margin, increasing contrast. Adjusting the full RGB (or CYMK) spectrum is often the most useful approach, but additional options allow color channels to be adjusted individually.

2. The LAYER function allows transformations to be stored in separate transparent overlays which can be added, removed, or altered in opacity. Results can be switched on and off or re-adjusted at a later stage. In the "virtual restoration" of damaged music leaves, good results have been obtained by selectively combining layers and by adjusting layer-opacity between 1 and 100%; a dark color response can be cancelled out by superimposing its negative counterpart, or color can be "multiplied." Sequential overlays can also be shuffled to achieve optimum results. Once an ideal result is achieved the overlays can be permanently "flattened" into a single image, saved separately from the original.

Layers can be imported from other documents and superimposed in the same way as those created within the document. This has proved an essential tool in incorporating the benefits of ultraviolet photography into "virtual restoration." Ultraviolet photography can be used to reveal faded material, but data visible under natural light may disappear. An ultraviolet capture can be imported as a partially transparent layer, adding new information without obscuring that unique to the standard capture under natural light: the discovery of a new composition in the Worcester fragments (Worcester Cathedral, Add. MS 68) was the direct result of this technique. Multiple major alterations made to a document can be saved separately and then re-introduced as semi-transparent layers to the original. A superimposed layer may improve one section of a document but obscure another; the interactive fluidity of adding, removing and adjusting layers is thus an essential part of both the process and its result. *Photoshop* allows an image to be saved (in its proprietary format) with layers intact, and users may themselves further adjust layer visibility. This is a powerful feature but increases file-size in direct proportion to the number of layers saved.

3. The COLOR-RANGE selection tool enables the user to select pixels of a specific color or color range. Up to 200 adjacent color-values may be selected (using a "fuzziness" slider), allowing the often uneven range of colors in, for

example, a damaged area of the page to be selected as a group and then to be replaced by a neutral background tone or otherwise manipulated. This tool has been used to lift layers of dirt, glue, water-staining, and even later over-writing from damaged pages, most spectacularly (to date) in the case of the palimpsest pages in Oxford, Corpus Christi College, MS 144. The best results frequently entail repetition of the process, since there is rarely sufficient contrast between "damage" and original ink colors. Independent range selections may be compiled in one action and colors selected more accurately by enlarging the image to a point where it pixelates. The effectiveness of these techniques is highly dependent on color bit-depth. Most capture software allows the camera operator to sharpen an image at capture, using "unsharp mask" or a similar tool. An unsharp mask increases the contrast gradient between pixels of a different color. At capture, this falsifies image data, producing an image that contains pixels of a color-value that do not exist in the original. But the effects can be spectacular when applied as part of the enhancement process, where the point at which the software recognizes a "difference" is defined by the operator, particularly in conjunction with other tools. Other filters may also enhance legibility. The 3D Transform rendering filter, for example, casts an artificial shadow, redefining light and dark as physical peaks and troughs. Dark writing on a pale ground appears embossed. Filters that expand dark pixels have also proved useful, to fill small gaps in written material left by the removal of other accretions in the "virtual restoration" process.

Addendum: Projects Exploring Similar Technologies

Other projects restoring data with digital images, using commercial software, and/or involved in the development of single-process applications include:

1. The Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology in Rochester, NY. Projects: the Dead Sea Scrolls, the Archimedes document and overwritten documents at St.Catherine's Monastery, Sinai.

2. Glenn Woodell: NASA/Langley Research Center Hampton, Virginia (Retinex Image Processing). Focus: Image-processing in various disciplines.

3. The Centre for the Study of Ancient Documents (CSAD), Oxford: recovery of ancient texts from papyri and tablets; development of repeatable algorithms for dealing with three-dimensional sources.

4. RCS: Computer science in the service of rock art research, protection, and education. For further information see *www.tpwd.state.tx.us/news/990913c.htm* and *www.infomagic.com/~rockart.*

5. Jim Henderson, RBP Applied Photographic Research (*hendersonjw@email.msn.com*): Henderson Cross-Polarized Enhancement Technique applied to American Indian pictographs and an ostracon from Qumran.

6. Nebraska State Historical Society (Assistant Curator of Photographs, Jill Koelling): restoration of 19th-century photographic negative plates.
Links to all these projects are available through *www.diamm.ac.uk/*.

References

Bent, Margaret (1970). *John Dunstable: Complete Works*, ed. M. Bukofzer, 2nd edn., rev. M. Bent, I. Bent, and B. Trowell, *Musica Britannica*, 8 (London), 212–214.

Bent, Margaret (1984). "The Progeny of Old Hall: More Leaves from a Royal English Choirbook," *Gordon Athol Anderson (1929-1981) in Memoriam*, Musicological Studies 49, 2 vols. (Henryville, Ottawa, and Binningen) i, 1–54.

Benton, John F., Allan R. Gillespie, and James M. Soha (1979). "Digital Image-Processing Applied to the Photography of Manuscripts, with Examples Drawn from the Pincus MS of Arnald de Villanova," *Scriptorium*, 33, 40–55.

Wathey, Andrew (1988). "Lost Books of Polyphony in Medieval England: A List to 1500," *Royal Musical Association: Research Chronicle*, 21, 1–19.

Wathey, Andrew (1993). "Manuscripts of Polyphonic Music: The British Isles 1100–1400," *Répertoire International des Sources Musicales*, Ser. BIV1-2, Suppl. 1 (Munich: Henle Verlag), 9*.