DIAMM

Digital Image Archive of Medieval Music

Digital Restoration Workbook

Produced in conjunction with the AHRC ICT Methods Network

Julia Craig-McFeely

Alan Lock

June 2006

www.diamm.ac.uk

www.methodsnetwork.ac.uk

An imprint of
Oxford Select Specialist Catalogue Publications

ISBN pending

prepared under the auspices of the
Arts and Humanities Research Council
Information and Communication Technology Methods Network
www.methodsnetwork.ac.uk

and the **Digital Image Archive of Medieval Music**
www.diamm.ac.uk

for the ICT Methods Network Workshop on Digital Restoration
29 June 2006

OSSC PUBLICATIONS 41 Freelands Road, Oxford OX4 4BS
Printed by the Alden Press, Osney Mead, Oxford

Contents

Workbook conventions

RGB refers to combined Red, Green and Blue channels

R, G, B means separate channels

R/G/B means either of the above (i.e. RGB or R, G, B)

To simplify matters the modifier key conventions are for the Windows platform, on the Mac substitute:

`Cmd` for `Ctrl` (`Cmd` is the designation for the `apple` key on older keyboards)

`Option` for `Alt` (newer keyboards already use the `Alt` key labelling)

Instructions for using tools or changing settings will be given both for Mac and PC, with the PC instructions first in each case.

Letter-keystrokes are always give in capitals so `Cmd-I`, never `Cmd-i`; `Cmd-L`, never `Cmd-l`

Menus and sub menus: when describing the location of a menu topic the convention will be to separate items using the ‘>’ symbol: e.g. `Image > Mode > Assign Profile`…

Location of files: when describing the location of a file or system resource the convention will be to list the top-level directory first with subsequent directories separated by an oblique stroke (‘/’): e.g. System Preferences/Desktop & Screen Saver/Desktop

Filenames referring to files used for the Restoration workshop are given in < >.

This workbook was written for a digital restoration workshop organised by the AHRC ICT Methods Network, and is therefore designed to be used in a workshop environment where questions raised by the content can be addressed by the tutors. However it should be possible to understand and learn the digital restoration process from the book alone. The main difficulty in doing so will be the absence of the images on which the workshop and therefore the workbook was based, and which the participants had to hand in order to learn and try out the techniques for themselves.

The software required is Adobe Photoshop, and screenshots of menus etc. are taken from version 7 of Photoshop on both Mac and PC platforms. Although some of the menu design is different, the content is the same.

In order to obtain images of a suitable quality for digital restoration, with which this workbook could be used, please see the end of Chapter 1 for image specifications.

1

SPECIFICATIONS

Computer Requirements

The sort of images necessary for digital imaging result in very large file-sizes – from 80 MB to 260 MB. Once the process of digital restoration begins that file size grows, since much of the work involves creating added layers, which increase file-size.

The first requirement is obviously a good deal of free disk space on the hard drive running the image-processing software: the image is best worked from the local drive to avoid speed restrictions imposed by communications protocol between the main drive and an external storage device. You may wish to store several copies of each image in different stages of restoration, which again multiplies the space required. Some machines run progressively slower as their hard disc space is filled up.

If your main file storage is network based you should make local copies of the image(s) you are currently working on. This is usually not only a lot faster for you but also a lot friendlier to other network users. Of course once the edits are finished the files can then be copied back to the network store to free up local disc space – if the number/size of files is large then this can be done when the network is less busy (e.g. overnight(.

The optimum working situation is to have a computer with the fastest processor available (or at the very least Pentium 2/Apple G3) and 1 GB of RAM or more. It is possible to manage with 512 KB of RAM, but you will find many processes are slow, and navigating the image is also tedious. In this situation (and to save time during the workshop) working with small segments of images may be advisable.

Customizing workspace

Note: Here the term 'workspace' applies to the computer and its environment, Photoshop has its own workspace options – this includes things such as menu layouts, colour settings and similar matters. These workspaces can be saved, recalled and even shared between users on other computers. This makes it easy to change Photoshop configurations on the fly according to the nature of the work being done.

It cannot be emphasised strongly enough that any computer used for quality image work must be properly set-up and calibrated – this includes any software used for the same purpose and the room in which the processing will be carried out. The human eye is a very flexible device and it will tend to see as white the dominant lightest area, this is the main reason we need to work a consistent, neutral environment. A simple checklist:

- try to work in a room that does not suffer sudden lighting changes
- try to equip the room with appropriate lighting (that rules out most fluorescent lamps – including 'energy saver' types)
- if possible ensure the room does not have large blocks of strong colours
- try not to compromise on monitors/display card quality
- properly set-up and calibrate the display physical properties (geometry, colour temperature, focus etc.)
- set-up your operating system for optimum display performance (drivers, refresh rate, 24/32 bit colour etc.)
- set a mid-grey background colour for your desktop (see later)
- work at an appropriate resolution (i.e. the addressable resolution) for your monitor
- take regular 'eye-breaks'
- if you wear spectacles do not use tinted/photochromic lenses for imaging work
- give your eyes time to adjust to the lighting of your imaging work area (particularly when coming from an area of very bright lighting) – a useful time to get on with some other type of work

> ☝ Throughout the digitisation, processing and delivery of your images do not compensate for users with improperly set up equipment – keep your systems running optimally and educate end-users, otherwise you run the risk of alienating those who have gone to the trouble of getting/setting their equipment correctly as well as compromising the quality of your archive.

Ambient room light

Part of the above is getting the room lighting right. Ideally the more critical the accuracy of your restoration the more careful you need to be about all of this. Ideally the ambient lighting around your monitor should be around 50% of the maximum brightness of your display (i.e. the white areas). This is so that the mage can be both brighter and darker than its surroundings, e.g. if the lighting was very bright then the screen could not match it and images would tend to look very washed out, the chances are that you would unconciously try and compensate for that, ending up with images that would appear saturated/contrasty under good lighting conditions.

To ensure correct colour evaluation it is always best to have a gray background for your screen, so dispense with any custom photographic, or brightly coloured desktop patterns. Monitor settings are changed as follows:

> PC: Start/Settings/Control Panel/Display/Appearance and choose the Desktop colour from there (e.g. 50% grey = R128, G128,B128)

☝ A short cut to the Display Control Panel applet is to right-click on the desktop and choose Properties – naturally this applies to the following two settings as well.

> Mac: System Preferences/Desktop & Screensaver/Desktop/Solid Colors/ select: Solid Gray Dark

Monitor Settings

Most monitors come out of the box with poor colour settings which are deliberately bright and over-saturated to make everything look more interesting. Older PCs often came out of the box set to use only 256 colours in the display: this was done to optimize performance, as it is much faster to redraw a screen with this setting, however it causes images to appear grainy and unclear. This needs to be changed before you can do image-processing work. Macs come out of the box with the screen set to millions of colours, but instructions for changing both type of screen are given below:

> PC: Start/Settings/Control Panel/Display/Settings/True Color (32bit)
>
> Mac: System Preferences/Displays/Display/ from drop-down menu on the right [`Colors`:] select `Millions`

Monitors also do not display at the best quality if the resolution set ofr their use is different from the optimum display settings of the monitor you are using. Generally this affects PCs more than Macs, but is worth checking as well.

> PC: Start/Settings/Control Panel/Display/Settings/Screen Area (set to native resolution of your display for flat panels or refer to monitor manual for CRTs). N.B. For PCs using CRTs it is usually best to set the highest refresh rate that your display/card can handle – it is better to drop the resolution slightly than run at refresh rates less than 75Hz.
>
> Mac: System Preferences/Displays/Display/ from the list of Resolutions choose the one that matches the screen size you are using. In some system preferences there is a check-box options to show only resolutions that suit the monitor. However, because newer Macs handle display settings better this may not be relevant. If you are not sure whether you are set to the right resolution, try changing between a few of those offered and see if the display improves.

Having changed these settings, there are simple adjustments you can make to make your monitor display colours more accurately.

> Go to this URL:
>
> http://www.diamm.ac.uk/content/access/check.html

The image on the page should show a smooth gradation of colours from white to grey to black. Check that the white is a clean white by comparing with a piece of white paper beside your monitor. The edges of the text and colours should be smooth if you have correctly changed your monitor to high-resolution settings. If the image you see does not match in any of these particulars (if it is not clear, sharp and lacking in graininess, you may need to check your colour resolution settings again), you will need to perform basic calibration on your monitor. It may be advisable in any case to check your monitor settings, even if you think the settings are correct.

PC:

Follow the procedure for setting colour depth above if you have not already done so. Then click the Advanced button in the Settings tab. Choose Color Management and Add or Set as default the profile for your monitor. If this can't be found then:

- Check for a driver disc that might have come with your monitor or...
- Go to the manufacturers website and download the driver
- If the either of the above isn't successful then set the profile
    - Go to the Windows `Start` menu
    - Select `Settings`
    - Select `Control Panel`
    - Double click on the `Display` icon
    - Click on the `Settings` tab
    - In the `Colours` portion of the dialog box, select `High colour (16 bit)` or `Millions of colours`, then click on the `Advanced` button at the bottom of the dialog box.
    - Select the colour profile that matches your monitor. If your monitor description is not available choose sRGB Colour Space Profile as this is a generic colour profile.
    - When you have completed these steps, Click on `Apply` and `OK` to exit the control panels. The on the screen should now show a smooth gradation of colour from white to black, and ideally you should be able to differentiate between A, B & C, and X Y & Z.

Mac:

Classic OS

- Select the `Apple Menu`
- Select `Control Panels` (this can also be found from the desktop in your System Folder)
- go to `Monitors & Sound` (or `Monitors` on older systems)
- Click on the Monitor symbol; in the `colour depth` dialog box below, select `Millions of Colours`
- Click on the `Colour` symbol; click the `calibrate` button
- Follow the steps described in the control panel. When determining the gamma, it may help to squint or sit back from the monitor until the images on the screen are blurred: this will enable you to see when the solid symbol blends into the background. The white point setting should be 6500. This will look very yellow/pink at first. switch to 5000 and then back to 6500, then continue.
- When you have completed these steps, close the control panel. The image on the screen should now show a smooth gradation of colour from white to black, and ideally you should be able to differentiate between A, B & C, and X Y & Z.

OS X

- Click on `System Preferences` in the dock or on the Apple Menu
- Click on `Displays`; from the drop down `colour` menu on the right select `Millions` - this should already be selected as OSX computers use high-bit colour resolution as default.

- Click on the `Colour` tab at the top of the dialog box; Click on the `calibrate` button.
- Follow the instructions for calibrating. On the first pane, click in the `Expert Mode` button to turn on additional options.
- After setting the target colours, set `target gamma` at 1.8
- Set D65 for `target white point`
- Complete the remaining steps by giving the profile a name. The image on the screen should now show a smooth gradation of colour from white to black, and ideally you should be able to differentiate between A, B & C, and X Y & Z.

If you have image-processing software, you can achieve better colour reproduction on screen by going through the steps described in the software to colour calibrate your monitor more accurately.

Note that Photoshop comes with the Adobe Gamma Utility and that this can be used to create (or modify an existing) profile. If you haven't used it before then running through the process is very instructional. Unless you are using a calibration device (see later) it is very well worth running through this process even if you already have a profile for your monitor. The reason for this is that the supplied profile are for *average* devices of that model: the age of the device and the conditions that it is used under will mean that the profile, although a good starting point, is unlikely to be optimized for your circumstances. Another advantage is that unless the lighting conditions are extremely consistent it is more accurate to have a set of profiles for the range of conditions under which the monitor will be used. Changing the current profile is as easy as picking it from the list (as outlined above) and applying it.

Colour calibration

What is it and why is it used?

Correct calibration is not essential for image restoration, since you will be changing the colour and appearance of the image you are working on. However understanding what is required and involved may be important when ordering images for restoration work.

Put at its simplest, calibration is a set of information parameters that tell your hardware and software how the colours on the image behave. A colour profile is needed for every item in your workflow that touches the images:

- the capture device (camera, slide scanner, flatbed scanner)
- The monitor used to view the images
- The output device used to print the images (inkjet, laser etc.)

Even 'duplicate' items bought at the same time, from the same manufacturer do not react the same way to colours: two high quality digital capture devices will scan the same basic picture and produce slightly different colour results.

Similarly, duplicate monitors will not have the same colour responses across the spectrum, nor will matching printers output the same colours as each other.

We have a tendency to assume that because we are working in an electronic medium, where things are either 'on' or 'off' that once calibrated, this calibration does not change, but it does. Even high-end equipment shows colour drift over time as components age.

In order to capture pictures that will look as close to the original as possible, no matter how they are viewed or output, the device that scans or photographs the images must be correctly calibrated

Calibration software: how it works.

First of all a calibration target, which contains specific predefined colour information is scanned or photographed.

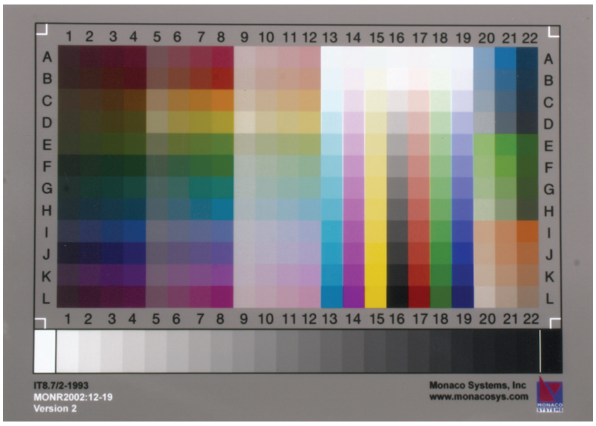

Sample Calibration target as it should appear

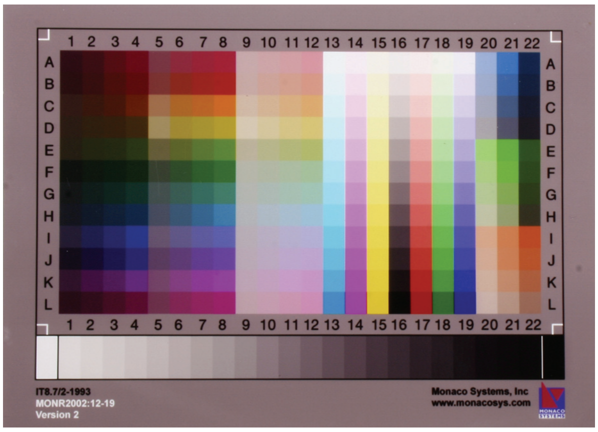

The same image with the wrong colour profile applied (i.e. one for a different capture device)

Then the resulting digital image is 'shown' to calibration software which compares the colours in the new image to the colours it expects to see. It is able to recognise very slight colour variances and it creates a profile based on the discrepancy between the way the calibration image *should* look and the way it actually *does* look. When the calibration profile is applied to the test image there are subtle shifts in colour which may or may not be visible to the eye.

However, this image may now look 'wrong' on the screen, unless the monitor is calibrated too. For this purpose, calibration software has electronically stored colours in its memory and it knows exactly how these colours should look. A device called a 'spider' must be attached to the screen over a patch which the calibration software creates. The spider reads the colours and recognises if there are any discrepancies between what it expects to see and what it does. Just as with the image, it then creates a profile that adjusts the appearance of the colours on the monitor so that the display on the screen looks correct. After calibration many monitor colours look dull and under-saturated, as 90% of monitors are sold to display colours bright and vibrant, to make the impact more attractive.

Once your monitor is calibrated and you can view an image with the correct calibration profile embedded in it (so the software knows how to display it) what you see on the screen will be close to correct colour. Unfortunately in all but specialist monitors, you may not be able to get exactly close to the correct colour.

The most important thing is never to adjust the image by judging colour from the on-screen display. Even if the calibration is very accurate you are nevertheless viewing an image with transmissive light, when the original only reaches your eye through reflected light.

You do not need a correctly calibrated screen in order to print images accurately, which is why calibration software is sold both with and without the screen calibration part.

In order to output your image with accurate colour reproduction, the calibration software also needs to create a profile for the colour behaviour of your printer. A preset image from the software is printed, but this must then be scanned so that the software can 'read' the output result and correct for any incorrect colour output. It is impossible to calibrate a printer without some way of getting the printed image back into electronic form.

If you use different weights, makes and types of paper (e.g. ordinary paper, coated inkjet, gloss-finish photo paper, high-quality matt photo paper etc.) you will need to create a profile for each make and type of paper. Many printers come with software that includes profiles, but these are ‘generic’ and although they will produce a good colour reproduction, it will not be accurate.

A purist will recalibrate for every new box or packet of paper that opened, but this can be costly, and the variance is only slight. However recalibration at regular intervals is advisable, as the printer components age, and is essential to recalibrate every time you change your ink cartridges.

Image specifications

Resolution; Image sizes and dpi

The way this term is used depends on what is being discussed. It is often used to refer to image size (in pixels), the resolution (in pixels per inch, ppi) of the image acquisition device (e.g. scanner) or the dots per inch (dpi) information within an image file that will be used for scaling printed output.

For example: if we scan a 3x2” image at 600 ppi we will have an image size of 1800x1200 pixels (at 600 ppi), if we send this to a printer at 300 dpi (ignoring the 600 ppi setting in the file) we will end up with a 6x4” print. If we use the 600ppi resolution information that should have been embedded when the image was scanned when printing the image then we will end up with a 3x2” print i.e. the same size as the original.

Notice that in this example there are direct correspondences between the size of the original image size, scan resolution and print size. However for digital cameras this is not the case: the resolution in ppi is a function of the size of the camera sensor, focal length of the lens and the distance to the object being photographed – this is the reason why we need an object of known size in the image to provide scale information.

The greater the number of pixels in an image the larger the file size will be (we are talking uncompressed files here). What is often unappreciated is that the files size increase as the square of the increase in resolution e.g. an 800x600 pixel RGB image will take up 1.44MB, if we double the resolution to 1600x1200 then the storage size goes up to 5.7MB and if we double it again it will go up to almost 23MB. These figures are for 24 bit RGB files, if we increase the bit depth to 48 bits (see later) then the storage requiremenst double again (to 45Mb in the last example).

File formats

What JPG or ‘lossy’ compression does to an image:

JPG/JPEG: Joint Photographic Expert Group. There are two main types of this format (which obviously tends to cause confusion): the ‘original’ JPEG and the newer JPEG 2000. The original was designed for maximal image compression. It uses a kind of compression called "lossy" compression – information is lost but in such a way that, when the image is decompressed, the human eye won't find the loss too obvious (some refer to this as ‘visually lossless’). The important point to note here that image information is *permanently* discarded, it is irretrievable. The amount of compression is variable and the extent to which an image may be compressed without too much degradation depends partly on the image and partly on its use. JPEG images may also degrade further if manipulated further and saved. Make sure that if you have to compress your work, compressing the image as a JPEG image happens last, after all manipulation or keep the original as a lossless format. It is much used as format for sending 24 bit colour images where the file sizes, even in a format using lossless compression, would otherwise be too large.

The JPEG 2000 format can use both lossless and lossy compression. Naturally the lossless will not lead to degradation of the image data, the lossy compression is based on wavelets and can give higher quality images for the same compression ratio as the original JPEG (or higher compression at the same image quality). This format is available in the later version of Photoshop (CS onwards, not in Photoshop 7) but is not installed by default.

> ☝ Photoshop does not allow the JPEG images to be compressed to anywhere near the extent of some other image editing software, this leads people to the erroneous conclusion that the JPEG compression algorithms are better in Photoshop.

The following images are created from the same base image file.The first is saved as a TIF, without any compression:

TIF

In a compression algorhythm the software decides that two adjacent colours which are nearly similar are similar enough that they can be recorded as the same colour, thus saving storage space. In the picture above this has caused the smooth gradation of greys to break up into blocks.

The next two versions of the image were saved in another program that, unlike Photoshop, allows the full level of JPEG compression to be applied.

This program uses a scale of 1-99 to represent the amount of compression, the first image (JPG1)was compressed to a level of 80 and the second (JPG2) 99.

N.B. The effect of the compression is not seen until the file is reloaded, something that has caused some users quite a bit of grief.

The most immediately noticeable change in the second image is that the whole thing is paler. On closer examination the cloud formations on the right have mostly been lost, but more noticeable is the breaking up of the smooth tonal change across the beach and the large cloud on the left. The extreme effect of this compression on the third image is quite startling.


JPG1


JPG2

Bit depth

This term is used to refer to the (maximum) number of colours that can be represented in an image file. The scale is binary based: 1 bit would give two colours (like faxes), 2 bits 4 colours, 3 bits 8 colours and so on. The most common bit depths are 8, 12 and 16 bit for greyscale, representing 256, 4096 and 65,536 shades of grey respectively. For RGB images there are the same number of bits per R, G, B channels i.e. 24, 36 and 48 bits giving 16.7 million, 69,000 million and 280 billion colours. Colour images can be saved (e.g. GIF files) at lower bit depths such 8 *total* for R, G and B – this gives 256 distinct colours but are not relevant here.

Notes:

The most common file bit depths are 8 and 16 bits per channel so a 12 bit capture would need to saved in a 16 bit file if we wanted to keep all the tonal information.

These above numbers of colours/greyscale refer to images which make full use of the available bits in the capture device, if the exposure is one stop under then an 8 bit image effectively becomes a 7 bit one (with 128 grey/8 million colour tones max.) Two stops under halves this again and so on.

Having pointed out the restriction that underexposure can lead to it is *imperative* that an image is not overexposed, unlike film there is no chance of recovering lost highlight information – it simply isn't there. This is why, if at all possible, it's best to save the image data at the maximum capture depth of your device: an underexposed 12 bit image can be downsampled to a full range 8 bit image. The only downsides to this are: greater storage space requirements; heavier demands on RAM; longer loading and slower processing times – and some imaging tools don't like high bit depth images.

☝ It is worth noting that reducing colour depth is a form of lossy compression – the tonal data lost is not recoverable (except by going back to the original file that you have kept safe!

RGB vs CMYK

CMYK (Cyan, Magenta, Yellow, blacK) is really only used for print output, but even here most normal desktop printers take care of the conversion from RGB to CMYK. Very few capture devices generate CMYK directly, the gamut (range of colours it can represent) is smaller, the files are 1/3 as large again compared to RGB and a number of imaging tools only work on RGB files. So, in essence, unless you have good reason to work in this space (and you'll know if you have) it's best to avoid it.

Post-processing

A maxim in image-creation is that if you have to go back to the image after taking it to adjust it to 'correct' it, then the image was not taken correctly in the first place.

There are few exceptions to this rule: specifically, if a book cannot be opened 90° or more, then it is almost impossible to get the camera perpendicular to the page, and a rectangular page may turn out rhomboid in shape. This may also cause the focus to be uneven across the page.

However, most post-processing is done to correct deficiencies in the original capture process.

Level or contrast adjust is sometimes done to correct poor exposure; colour adjustments may be done to correct poor grey-balancing of the capture equipment, to compensate for lack of (or incorrect) colour calibration, or even to correct the screen-viewed image to make it look the same as the original. All of these activities destroy colour information that is vital for image restoration, and will often even falsify colour information, so that any possibility of restoration is compromised.

De-skewing (straightening up of crooked images, or squaring up of rhomboid images) is sometimes done to correct poor placement but, no matter how irritiating it may be to see an image slightly crooked in your viewer, that is better than having separate colours blurred into each other by the de-skewing process.

An unsharp mask (in addition to its general role of adding 'punch' to an image just prior to output) is often applied to compensate for poor focus, either because the operator has failed to focus the camera tightly enough, or because the camera is too far away from the image: some capture software offers the option to include this at the point of capture, thus falsifying image information very early on in the process. Unsharp Masking or other sharpening applied after capture changes colour values, and is particularly destructive in the fine gradations of colour which form the part of the image that is most vital in enhancement processes. The image shown above has been sharpened: when enlarged on screen to actual-pixel view you can see that the smooth background colour has become grainy and broken up, thus losing the subtle colour definition crucial to restoration. Even with the sharpening process, the image is not clear: the edges of the red are not clearly defined, and the small brown mark between the letters is nearly impossible to see.

> Digital cameras (i.e. most consumer and prosumer cameras, the exception being those that use the Foveon sensor) that use array capture devices employ averaging techniques to obtain the full colour information for their native resolution. Inevitably this introduces softening of detail which is why some form of sharpening is required.

In order to apply digital restoration processes to an image, the quality of the original is crucial:

- It must be in focus to the smallest detail – hair follicles on parchment, or the grain of the paper on paper. This is not as demanding as it sounds, and all images should be this sharp as long as the source is not exceptionally large, in which case it should be shot in sections with the camera closer. Subjects that are not flat introduce their own problems, normally you would use a smaller aperture to increase depth of field but the longer exposure time required may rule this out.
- It must have as wide a colour spectrum as possible, so there must be **no post-processing adjustment** that might change the spectrum or values of colours captured.
- The file must be captured and delivered in an uncompressed format to avoid data loss through compression.

Data content

In order to evaluate the quality and colour of your image, you need some control data in the picture itself that will answer many questions about the image content that is not otherwise apparent.

When viewing images, the original artefact is almost never at hand. However a centimetre or inch is always the same, you just need to know the scale in relation to the object. Colour cannot be defined so easily, but we use scales that are industry standard: i.e. you can get one anywhere, and it should always look the same.

When we look at an image with a scale, we have simultaneous control information that tells us whether what we are looking at accurately represents the original artefact or, if not, in what way the information may differ from the original.

You have the size, you can also judge the exposure (light and dark) and the veracity of the colour reproduction, either in the image, or in the viewing medium, or both.

From the technical side, scales allow the photographer to calculate correct exposure and adjust for any colour cast caused by the environment, it also allows the final image to be evaluated for quality control.

Where the image is saved as a TIF, there is space in the document ‘header’ (not visible on the image) in which metadata about the image and the capture environment should be stored. This is a safety measure in case control data about the image is lost, or the filename becomes corrupt. It is also general good practice in creating digital images. The content of the TIF header can be seen in Photoshop by choosing menu: `File > File info…`.

Software and Saving

The processes described in this book are performed in Adobe Photoshop version 7. Although there is other good photo-processing software available, the specific tools provided by Photoshop make it particularly suitable for the type of image restoration described here. Photoshop CS/CS2 superseded Photoshop 7, with the notable change that CS is able to better handle 16-bit grayscale and 48-bit colour images i.e. many more tools are available for use with this type of image. Some features have changed marginally, but not to the extent that the contents of this workbook could not be applied using Photoshop CS.

Because of the size of the images, it is often not possible to open them in non-image-processing software. Cut-down versions of Photoshop or other software provided with digital cameras may open the images, but will not allow the level of editing required for restoration.

Images saved using advanced features of a particular image processing software may not be supported by other image applications – thay may either ignore the feature or refuse to open the image. This is one of the reasons that images for final delivery will usually need these features removed – more on this later.

Photoshop recognises most file formats, though Version 7 tends not to recognise RAW formats generated by digital cameras. These are read by CS.

In addition to the general-use file formats, Photoshop also has its own proprietary format, which enables you to save work utilising Photoshop layers. It is called simply 'Photoshop' format and the file suffix is .psd. When creating new copies of documents for restoration, save as .psd, since this enables you to save restoration layers so that they can be adjusted or removed in a later work session.

As with any work that you are doing, it cannot be emphasized too strongly how important it is to continually save your work. If the RAM cache overfills on the computer you may find that a long work session cannot be saved, because there is too much information in the memory buffer. It is also not exactly unknown for software to crash suddenly for no apparent reason, and you do not want to lose more than 5-10 mins work at most. Every time you pause for thought, or complete a task, `SAVE`! If you feel particularly paranoid you may want to save a series of copies of your file at different stages of work. File corruption is also not unknown.

Human specifications

Colour-blindness

A significant percentage of men (~7%) and a very small percentage of women (<1%) suffer from colour-blindness. A surprising number of people are not even aware they are colourblind. Although it is not essential to have perfect colour perception for image restoration processes, it can be extremely difficult to differentiate between the fine gradations of colour if you are colourblind. A quick online test will give you some idea if you have a colour perception deficiency, and a vision specialist should be consulted for more accurate testing (although there is no 'cure').

Online vision tests can be sampled here:

http://colorvisiontesting.com/online%20test.htm http://colorvisiontesting.com/ishihara.htm

For more information in understanding colour perception:

http://www.iamcal.com/toys/colors/

Working practices

Protecting your eyes

Digital restoration can be compulsive: don't sit doing this for hours on end, or you risk serious damage to your eyes. There are a number of 'resting' exercises you can do to minimize eye-strain such as looking at distant objects out of a window at regular intervals, or even just focussing at objects within the room which are at a different distance from your work screen. Try not to work with your monitor at its brightest setting – ambient room-light is particularly relevant to this factor in eye-strain management

Ordering images from Libraries or Archives

It is important to remember that the provision of photographs or digital images to researchers is a courtesy on the part of the library. Although most libraries will bend over backwards to provide researchers with what they need there will be times when you just can't get the quality or size of image that you want.

Nevertheless the images you get should meet certain quality standards, and if you receive a slide or other image that is e.g. out of focus do not feel unable to return it to the library and ask for a better copy.

High-resolution imaging equipment necessary for the creation of restoration-quality images is extremely expensive, and only a few major archives have access to this level of technology. However, the cost of imaging equipment is coming down steadily, and if you can wait you may find that the type of images you require become available within a couple of years. If you have specific needs which the library concerned cannot fulfil please contact DIAMM, and we may be able to help you obtain what you need.

Technically the copyright holder of the image (that is, the person who took the photograph or their employer) must give permission for you to manipulate or otherwise alter their image, but in practice this is not usually a problem unless you need to publish the results of your work, in which case you need permision to publish an altered version of the image, and you must take great care to indicate that this is not how the original image, or the original source appears.

The main concern of libraries is that their holdings are not misrepresented in any way. The ideal if you need to publish a manipulated image is to publish the original alongside it (appropriately captioned).

NB: putting an image or part of an image on a website constitutes publication, and infringes copyright unless you have permission to put the image online from the original creator and owner of the document photographed.

It is a matter of courtesy to ensure any images published or put online are accurately captioned as to their origin: you must include the country (our country abbreviation), city, name of the library, accurate shelf mark and exact folio of the image you reproduce. You would normally also need to acknowledge the copyright holder and sometimes the owner if these are not the same. Correct captioning saves libraries and scholars hours of time trying to locate materials that are incorrectly or incompletely captioned. It is also important to supply the library with a copy of any publication arising from use of their image(s).

Anything you do with digital images ordered from a library will have an effect on how the document owner responds to future requests for images or access to their document(s): digital images are very easily copied and passed around, and this infringes copyright and is a breach of the agreement you sign when you buy the image. Easy access to high-quality digital images in the future depends on careful observance of rights today, while access mechanisms and policies are still being formulated.

Check-list for ordering images

- Specify that you want High Resolution (and say 400 dpi at real size or higher);
- Uncompressed TIFF format both at capture and in delivery format; NO JPEGs in the workflow;
- 8-bit (or 24-bit depending on library protocols for describing bit-depth) RGB colour minimum (this is the standard setting, 16-bit (or 48-bit) creates much larger images but also stores more colour information. 48-bit images can only be opened using Photoshop CS);
- Colour profile of capture device embedded in the image;
- Industry-standard colour, grayscale and size scales photographed beside the image but not touching it;
- White and black points adjusted to avoid clipping at the point of capture (not afterwards)
- Image sharp (i.e in focus) to the finest level of detail;
- NO unsharp mask applied during capture;
- NO unsharp mask applied after capture;
- NO rotation, deskewing, reshaping, levels, colour or exposure adjustment after capture (i.e. no adjustment of any sort after capture)

☝ In the UK images of restoration quality usually cost in the region of £25-30. If you are intending to order a large number of images (e.g. a complete MS) it is worth contacting the library to ask if they will consider doing a special bulk price. In the USA and Europe you may pay a similar or lower price for images, but reproduction rights are 3-10 times the price of the image in addition to the purchase price, particularly reproduction online.