



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

© 2006 Julia Craig-McFeely & Alan Lock
Julia Craig-McFeely and Alan Lock reserve the right
to be identified as authors of this publication

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	5
1 Specifications	6
a. Computer requirements.....	6
• customizing workspace.....	7
• Monitor settings	8
• Colour calibration.....	10
b. Image specifications.....	12
• Resolution; Image sizes and dpi	12
• File formats	12
• Bit depth.....	14
• RGB vs CMYK.....	15
• Post-processing.....	15
• Data content.....	16
c. Software and saving.....	17
d. Human specifications	17
• Colour-blindness.....	17
e. Working practices	18
• Protecting your eyes	18
• Ordering images from libraries and archives	18
f. Check-list for ordering images	19
2 Photoshop.....	20
a. Basics:.....	21
• Tool settings.....	21
• Keystrokes	21
• Palettes	22
• Write-protecting source files.....	22
• Different ways of achieving the same result.....	23
b. Tools.....	23
• Zoom.....	24
• Crop	24
• Selecting	25
➤ marquee	25
➤ lasso.....	26
➤ magnetic lasso.....	26
➤ magic wand	26
➤ select colour range menu	27
➤ adding to/subtracting from selections.....	28
➤ inverting selections	29

	➤ expanding and contracting	29
	➤ feathering.....	29
	➤ smoothing	30
	• Straightening; Cutting and pasting; resizing canvas.....	31
	• Layers	32
	➤ Creating a new Layer	33
	➤ Adjustment layers	33
	➤ Layer order	34
3	Restoration 1 (Julia Craig-McFeely)	35
	• Ethical Process.....	35
	• Sequence of operations	36
	• Global level adjust	37
	• Select colour range.....	38
	• Single-area level adjust.....	41
	• Colour fills	43
	• Pattern fills	45
	• Combinations	46
	• Filters and Variations	50
	• Layer transparency	51
	• Common errors and problems.....	51
	• Variations on colour select.....	52
	• How far to go	53
	• Unethical process	53
4	Restoration 2 (Alan Lock)	55
	• Using levels to assess an image	55
	➤ White and black points	55
	• More on Layers	56
	➤ Layer modes	56
	➤ Curves tool	58
	➤ Sharing Layers between images	59
	• Channels	60
	➤ The different types of information	60
	➤ Viewing Channels	60
	➤ Using Channels	61
	• Masks.....	62
	• Unsharp masking.....	64

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; customizing workspace

- Monitor settings.....
- Colour Calibration.....

Image specifications

- Resolution; Image sizes and dpi
- File formats
- Bit depth.....
- RGB vs CMYK
- Post-processing
- Data content.....

Software and saving

Human specifications

- Colour-blindness

Working practices.....

- Protecting your eyes.....
- Ordering images from libraries and archives

Check-list for ordering images

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 image 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 unconsciously 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 off 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 requirementst 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:



In a compression algorithm 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 as 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 be 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 irritating 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 – they 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 permission 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.

2**PHOTOSHOP**

Basics:

- Tool settings
- Keystrokes
- Palettes
- Write-protecting source files
- Different ways of achieving the same result

Tools

- Zoom
- Crop
- Selecting
 - marquee
 - lasso
 - magnetic lasso
 - magic wand
 - select colour range menu
 - adding to/subtracting from selections
 - inverting selections
 - expanding and contracting
 - feathering
 - smoothing
- Straightening
- Cutting and pasting
- Resizing canvas
- Layers
 - Creating a new one
 - Adjustment layers
 - Layer order

Basics

Tool settings

Changing to precise tool cursors

Default settings in Photoshop causes the cursor to represent the selected tool rather than being a pointer or something more accurate. For restoration you will need to use a more precise method of pointing, and the settings on Photoshop need to be changed.

Menu: **Photoshop > Preferences > Display & Cursors...**

For both **Painting Cursors** and **Other Cursors**, select **Precise** and hit **OK**.

The History palette records the actions you do from the time you open the document, but its memory is limited unless you change the settings:

Menu: **Photoshop > Preferences > General...**

For **History States** enter a larger number, e.g. 30. If you have a large RAM allocation you can set this number higher, but doing so uses RAM, and you will need to balance the requirement for saved History with the specifications of your hardware.

 The history is wiped out when you close the document and cannot be retrieved.

Conversely, if you are low on RAM you can free some up by purging the history, menu: **Edit > Purge**, there you will find a choice of items to purge (Clipboard, History, Undo and All). The Clipboard option is useful for a big Copy/Paste operation.

Keystrokes

Using keystrokes can vastly speed up your working process, since you will spend less time positioning the mouse. It is also a good working practice to minimize use of the mouse where possible, as over-use, particularly when you are not used to extended mouse use, can lead to RSI: restoration inevitably involves a lot of mouse work.

 The highlighted button in any dialog box is activated by hitting the **enter** or **return** keys. The **cancel** button is activated by hitting the **esc** key

To reset the dialog box to its original settings without exiting it, press the **alt** key (PC/Mac) and the **Cancel** button will change to **Reset**.

There are useful shortcuts within Photoshop using 'modifier' keys to change the effects of tools, to add or subtract or to change button designations

PC & Mac

Alt reverses the action of the tool (e.g. from *zoom in* to *zoom out*, or from *add* to *subtract*)

Alt allows temporary subtraction when addition is the current mode for selection tools, **Shift** permits addition when subtraction is the current mode.

Shift allows you to add to selections, **shift-alt** to subtract from them

PC:

Ctrl-Z undoes last action, **Alt-Ctrl-Z** undoes last step.

Mac:

Cmd-Z undoes the previous action; **Cmd-alt-Z** steps back one history state at a time (the same effect as multiple undo)

For multiple undo actions use the History palette.

The other nice thing about learning the keystrokes is that some are often (but not always!) active during the use of another tool, a simple example: you have the Levels dialog box up and have made several adjustments but you want to examine both the detail and overall effects and using the Zoom tool would mean closing the Levels dialog and losing your adjustments – instead just use **Ctrl+** to zoom in, **Ctrl-** to zoom out, **Ctrl-0** to fit on screen and **Alt-Ctrl-0** to view actual pixels (i.e. 100%), the scroll bars are still active for panning OR you can press the space bar and the pointer changes to the Hand tool.

Palettes

When you first open Photoshop, the tools and palettes are laid out in a default workspace. On the left is the tools palette, and on the right are a number of other function palettes. To save space, each palette window may have two or three tool palettes on it, accessed with tabs at the top of the window. Each tool palette has its own drop-down function menu, found under the ► sign at the top right of the small window.

You can close any of these windows and reopen them by selecting them in the menu: **Windows > [tool palette]**. You can also change the default layout by dragging the windows, resizing them or dragging tabs between windows to create different work sets. You can also 'dock' any tab and its content in the grey dock area at the right-hand end of the options bar.

Write-protecting source files

The first action you should always do when preparing to work on an image is to protect your original. NEVER work on the original file, as you may need to refer back to it. It is also sometime useful (if you have enough screen space) to have the original file open at the same time as the restored file, so that you can keep track of whether you are actually improving or disimproving with your work.

You can write-protect your originals by keeping them on a non-rewritable CD or DVD, but this type of storage is unreliable, and you may find that your original image is lost. If the images are precious, and you have the only copy, then make sure you have made multiple archive and backup copies.

Before beginning work, copy your image to the local hard drive. To write protect this image:

PC: In Explorer: right-click on the file (or group of files), choose Properties and check the Read-only Attribute box om the General tab.

Mac: click on the image icon, hit **Cmd-I** or Menu: **File > Info**

In the dialog box that comes up, click in the **Locked** checkbox. This will prevent you from editing this document accidentally.

Different ways of achieving the same result

Photoshop is an extremely versatile and powerful image editor. From users point-of-view this means that it takes quite some time to understand the complexity of the program. Another aspect of this versatility is that there are often a number of ways of achieving the same, or very similar, result. Each person may well have their own favourite set of tools, this can arise through familiarity (they understand or have a feel for the results they can obtain with that tool), or that they have never used other tools which may give them the same (possibly better) result. There is often a tendency for users to stick with what they know – sometimes they believe (and tell others) that this is the *only* way to achieve a given result. Try not to get into this habit, by all means use the tools you know best but also try and keep an open mind and explore techniques that you have either never used before or may have dismissed, books and websites can be good sources of new ideas and methods.

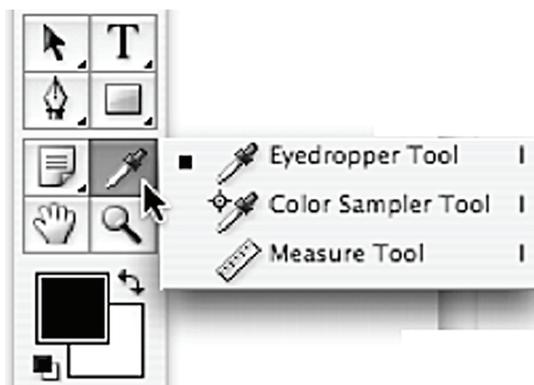
On the next page is a relatively trivial example of how a similar result can be achieved with different methods:

Tools



Photoshop provides the user with a palette of tools which more or less graphically represent what they are able to do. In order to avoid this palette becoming unmanageably large, related or refined versions of several tools are hidden under the primary tool icon. The tools with a little black tab in the corner have other tools hidden underneath, which show if you **right-click** on the tool (PC), or click and hold on the icon (PC/Mac), **alt-click** or **right-click** (Mac).

Occasionally a tool you have seen but cannot remember where to find is difficult to locate because it is hidden under something that is not related to its function. One example is the Measure tool, which is hidden under the eyedropper:



You can switch between tools by clicking on them, or by hitting the letter on the keyboard shown in the menu for each tool. You can also find out the keystroke for that tool by clicking on the tool and hovering over it until the popup help appears.

To cycle through the tool options using the keyboard on both the PC and the Mac hold down the **shift** key at the same time as typing the required letter (repeatedly).



Not all the tools have keystrokes assigned.



These images show how a similar effect can be achieved with different tools.

From top to bottom we have:

- Original,
- Curves adjust,
- Levels adjust and, finally
- Highlight & Shadow adjust.

The latter is available in Photoshop version CS and later.

There are obvious differences but the settings for the tools were not changed very much. With more 'fiddling' the results could have been made even more alike.



Zoom (keystroke Z)

This tool allows you to zoom your view in and out. There are tool settings on the options bar at the top of the display: you can change the default action from zoom in (with the + sign) to zoom out (with the - sign), but since you are likely to want to swap between the two it is best to leave it on one setting, and use the **alt** key to reverse the zoom direction. There is a checkbox which specifies whether the window resizes when you zoom, and also three buttons, which create specific views: fit on screen, print size and actual pixels. You will use this last button most, since restoration requires detail work at high magnification.

You can zoom on a specific area by clicking on a point in the image while the zoom tool is selected, but there is also a keystroke for zooming generally which allows you to zoom in or out not matter which tool you have selected: **Ctrl/Cmd +** or **-**



Crop (keystroke C)

To choose a particular part of an image and delete everything around it, click on this tool and then drag-select the area that you want to define. When you release the mouse button, a square or rectangular area is selected (to constrain to a square, use **shift** while dragging). The resulting area can be resized by dragging the corner or edge boxes. To see what you have selected more clearly, select **shield** on the tool options bar after dragging.

To deselect the area, click anywhere outside the selection, or hit **Esc**.

To delete the material outside the cropped area, hit **Return/Enter**.

Selection Tools



If you have created a selection using any of the tools described below but decide you don't want it, the keystroke to deselect is **Ctrl-D** (PC) or **Cmd-D** (Mac).



Marquee (keystroke M)

The Marquee creates specific shapes, squares, rectangles, circles, ovals. It is a click and drag tool and can be used in any direction. The corner of the square (with this shape selected) appears where you initiate your drag. From that point the tool creates a square or rectangle with the opposite corner sited where you release the button. Under this tool are circle and oval shapes.

To constrain the dragged shape to a square, hold down the **shift** key while dragging. Holding down **shift** while dragging with the ellipse shape constrain the selected area to a circle. When creating a circle the shape is created as if within an invisible square.

Although you can move the selection once you have made it by nudging (using the pointer tool) you cannot change its size except by adding or subtracting with another marquee, or by using a drop-down menu option.

☞ A common error is confusion between the Crop tool (which deletes everything outside the selected area) and the Marquee tool (which selects everything within the selected area). The Crop tool is resizable, but the Marquee tool is not. If you accidentally use the crop tool, simply use the document History to undo your action.



Lasso (keystroke L)

The Lasso is a tool that allows you to select any part of the image with complete freedom of shape. As you might expect from the name, you use the mouse (click and drag) to draw a shape and everything within the shape is selected. You will note that it has other tools hidden under this icon: these will be discussed later.

☞ When selecting with the lasso tool, the lasso shape will automatically join up to the starting point from wherever you lift your finger from the mouse button: you do not have to draw right back to the starting point yourself.

Magnetic lasso



The Lasso tool is a free-drawing tool, and can be difficult to manipulate with a mouse (a graphics tablet is significantly easier). Where you are trying to follow an irregular line such as the boundary of staining on a leaf, the Magnetic Lasso Tool is invaluable. It is found hiding under the Lasso tool on the tools palette.

This tool works by creating a selection path along a line based on colour contrast. The tool options bar offers options including 'Edge Contrast' where you can define the level of contrast that you want the tool to detect. To

create a selection path with this tool, click on a starting point and then drag the pointer along the boundary of the selection area. Selection points can be 'forced' by clicking, otherwise the software automatically puts them in at regular intervals. You can delete a selection point if you think it is wrong simply by hitting **Del**. To complete the selection area drag the pointer back to the starting point or origin of your path; the pointer will change shape and a click will join the two ends so that everything within the selection path is now selected.



Magic wand (keystroke W)

This tool allows you to choose a colour on the image and select all adjacent colours that are the same, or nearly the same. The degree of similarity is defined by the Tolerance number in the tool options bar. The higher the number, the wider the colour spectrum that will be selected. This is very handy for dealing with small areas or individual items on an image. You can add to the selection by holding down **shift** and selecting another colour or area, and subtract by clicking while holding down **alt**.



The icons in squares perform the following functions:



Selects the colours you click on: new selections cancel previous ones;



Adds to your previous selection;



Subtracts from your previous selection;



Selects a colour range within a previous selection;

Anti-aliasing essentially blurs the edges of things: so instead of something being a hard black and white edge – this is particularly noticeable on curves, which look jagged on black/white images. What anti-aliasing in this tool will do is fill in the missing bits on the curves with a mid-range colour to make curves look smoother. If you select 'Anti-aliased' the wand will create the fill-in colours on curves.

If you select 'Contiguous' this selects all the colours adjacent to the one you have selected, but not any similar or duplicate colours that are separated from your original point by another colour that is too dissimilar.

If you select 'Use All Layers' the colours selected will include all the layers you have created so far; otherwise it will only act on the layer which you have selected (check that you have click-selected the background layer if that is the one you are working on).

'Select Colour Range'

This is not on the tools palette, but is accessed from the **Select** Menu on the menu-bar. This tool behaves very similarly to the magic wand, but selects a much finer gradation of colours. It enables you to make fine-grained selections which select all the colours across the whole document, or within a preselected area. Another difference is that it creates a preview to show you what you have selected and, with large files, is much quicker to work with than the wand, since it does not actually select the pixels until you hit **OK**. It also provides some refinements in colour selections which may be useful. Before accessing this tool, zoom in on the image so that you can see individual pixels – 200% is good, higher if the area is difficult to see.



Before starting work using this tool, always make sure you have selected the background layer, or the one which you want your work to affect.

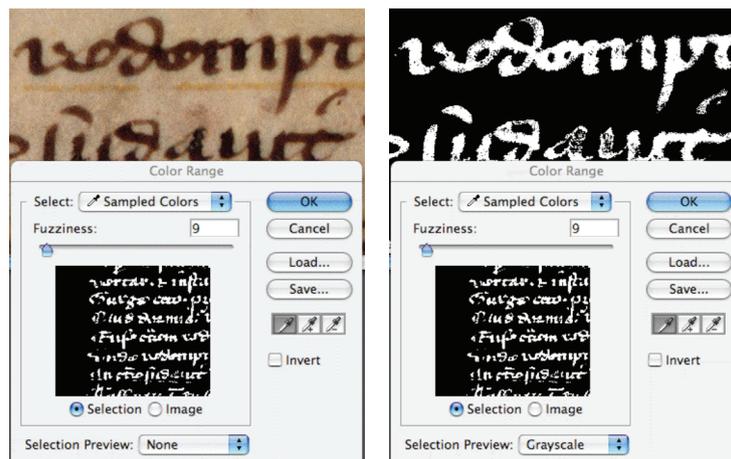
Menu: **Select > Color Range...**

The Dialog box that comes up will show a preview of your image in black and white. The white colours are the ones that are currently selected. In order to see the image in real colour in the viewer, hold down the **Ctrl** key (PC)/**Cmd** key (Mac)

If you click anywhere on the image behind the tool dialog box, the colour you click on will be selected in the viewer. If the image is not in the right position you can move it by using the scrollbars or roller on your mouse. The 'Fuzziness' slider performs the same function as 'Tolerance' in the Magic Wand tool: it defines how broadly the software interprets the colour you have selected.



The fuzziness affects all colours you have selected. With the magic Wand Tool the Tolerance only affects selections that you make after changing that number, so you can create selections with widely varied colour tolerances.



In this example the black colour of the text has been selected. Although you can only see part of the image when it is blown up on screen, the preview shows the whole image (albeit very small). The fuzziness is at 9 to avoid selecting anything more than the required colours. In order to see where the ink colour has not been selected choose one of the options from the 'Selection Preview' drop-down menu at the bottom of the dialog box. It is clear from the second view which pixels are missing from the selection and need to be added (mainly in the letters 'a' and 'u' in the lower line of text visible in this clipping).

To add to the selection shift-click, or select the '+' eyedropper in the dialog box. Although you can increase the selection by moving the fuzziness slider up/right (and decrease by moving it down/left) the refinement necessary in restoration generally requires a low fuzziness (in the range 8-24 usually), and a large number of selected colours.

☞ Fuzziness can be changed during selection by using the rollerball on the mouse (if you have one) or the up/down cursor keys.

Do not attempt to select more than one colour – e.g. red *and* blue, as the software will select all the colours of the spectrum in between, and that will be too broad a range.

Although not a difficult or complicated tool to use, this is crucial to many restoration processes. The skill is mainly in judging the balance between selection and fuzziness. With practice you will find that your colour selections become much more accurate and you can work more quickly: at first the process is slow, and can be frustrating.

☞ One of the most frustrating errors encountered in using this tool is to create a long and painstaking selection, then on hitting OK to find that almost everything on the image is selected. This is because you may not have returned to the background layer before beginning your selection process.

Adding to/subtracting from selections

You can add to (or subtract from) selections with any of the selection tools: a selection made with the magic wand can be added to with the Magic Wand, Marquee or Lasso, but not with the 'Select Color Range' tool (this will only select colours within the selection area you have already defined).

To add to a selection, click on the desired tool and drag-select (point-select with the Magic Wand) while holding down the **shift** key. To remove areas from the selection click on the desired tool and drag-select (point-select with the Magic Wand) while holding down the **alt** key.

Inverting selections

A useful process with areas that seem to defy colour-range selection using the Magic Wand or Select Colour Range tool is to try selecting all the colours that you know you *do not* want (instead of those you do), then invert your selection: this reverses your selection and selects all the colours previously excluded. Menu: **Select > Inverse** or keystroke **Ctrl-shift-I** (PC)/**Cmd-Shift-I** (Mac).

Expanding and contracting

Having created a selection you may find that it is too grainy for your requirements. You can expand what you have selected in tiny increments by using the menu: **Select > Modify > Expand...** You can choose to expand by any number of pixels: the most useful course is to expand by one pixel. This often joins up single pixels that are just separated and creates a more effective selection area. Likewise, you can contract your selection, menu: **Select > Modify > Contract...**

Feathering

Feathering is useful when you do not want to create a hard-edged selection: it softens the edges of a selection. When selecting (for instance) an ink of a particular colour, you will find on zooming in that the bleeds at the edge of the pen-stroke are often quite different in colour. If you select the main ink colour and then remove it, you are left with the edges of the ink like a ghost image. If you expand the selection by e.g. 2 pixels, and then feather also by 2 pixels, whatever fill you put in place has a softer edge, and 'removes' the writing more effectively.

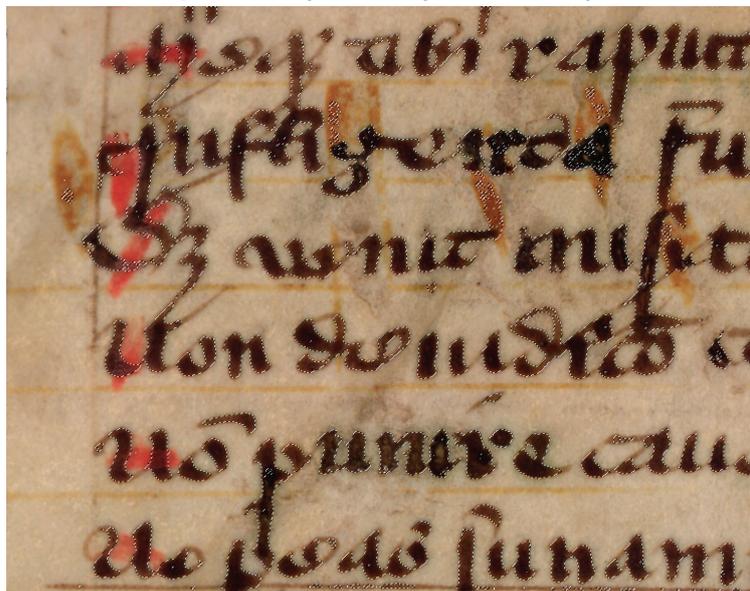
You can only feather when you have a part of the document selected. After creating a selection go to menu: **Select > Feather...** (note keystroke option on the menu) then enter a number for the pixel radius: the larger the number the wider the feathering will be.

In the examples below the over-writing (which we want to eliminate) is the black/dark brown text, and the writing we wish to see is mid brown/yellowish musical notation (diamonds and squares in this sample). The background is a paler pinky-brown. The first action was to create a pattern fill that matched the background (how to do this is discussed later), then a colour range was selected using the 'Select Colour Range' tool.

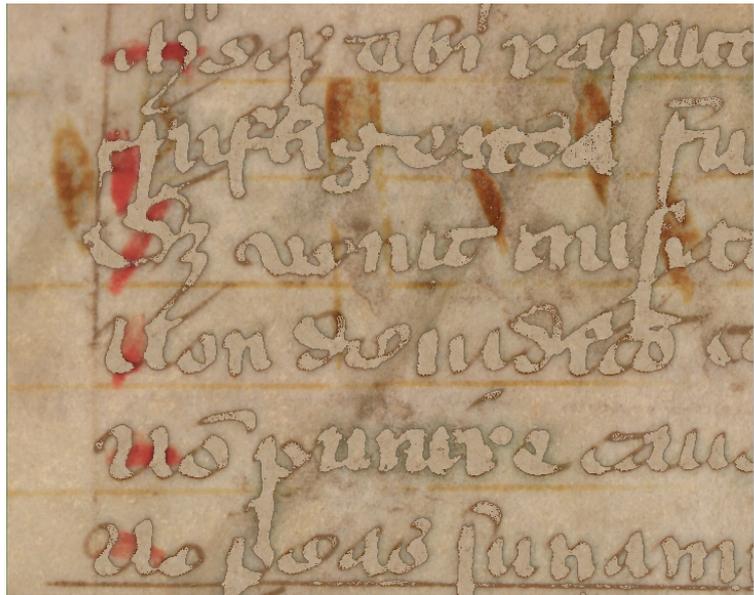
The colour range of the overwritten text was not particularly carefully selected – you can see that the selection is extending into the paler brown of the music notes – but even with a fairly generalized colour selection the difference in effect with and without feathering of the edges is dramatic. The softer edges of the filled area mean that they do not interrupt the eye to such an extent, and the palimpsest is revealed.

Steps:

- 1 Select ink colour using 'Select Colour Range'. (The edges of this image show the selection defined by a dotted line.



- 2 Replace selected colour with previously defined pattern based on background texture (described in Chapter 3). The writing still interrupts the eye when trying to read the underlying music notes.



- 3 Go back to previous colour selection; expand and feather selection by 2 pixels and then replace with previously defined pattern based on background texture.



This gives a result that interrupts the eye far less than the result after action no. 2. There would probably be a value in performing a further colour range select on the grey edge-colour, but this is sufficient to demonstrate the difference in result.

If you feel that your selection is too broad you can feather without expanding, or contract the selection using **Select > Modify > Contract...**

Smoothing

When a selection range appears very bitty and broken up, but expanding the selection may include pixels which you do not want to select, smoothing can have a more desirable result, joining up pixels at a proximity that you define in the dialog. Smoothing is found under the menu: **Select > Modify > Smooth...**

Performing any of these actions as a test is possible, since if the result is not as you had hoped, you can step back to your original selection using the Document History or Undo.

Straightening; Cutting and pasting; Resizing canvas

Straightening is inadvisable prior to digital restoration, unless you are rotating through 90° or 180°. Slight rotations are achieved without apparent visible effect on the picture, but in fact it does mean the rewriting of colour values and slight blurring of boundaries. As you can see from the example above, the boundaries of letters and artefacts on an image are crucial to defining content.

In the example below, the manuscript was cut in two, and the restorer wishes to reunite the two halves of the page using the digital images.



Steps:

Locate a control area on each page that will be used to straighten the images that will be the same between the two images. The edges of the page are not good control areas, as they are rarely straight, even at the point where the pages were cut in two. In this case the stave lines of the music or the vertical side rules are probably the best point from which to define vertical or horizontal. There are usually rules of some sort on pages of text which can be used in this way.

The tool used for straightening is the 'Measure Tool', found under the 'eyedropper' (keystroke I). Select the tool then draw a control line (using click-and-drag) along the line that you have chosen to be horizontal or vertical (in this case it will be horizontal). Having drawn the line exactly following the angle of the stave line, go to menu: **Image > Rotate Canvas > Arbitrary...**

The dialog box that opens has a number already inserted into the angle box – do not change this: it has been calculated by the software based on the control line that you drew with the measure tool, and is the correct angle of rotation required to turn this line to horizontal (or vertical, whichever is closest). Hit **OK**

The image will rotate and become slightly larger to accommodate the angle. Perform this action on both pictures that you want to join. In the image of the top of the page, you need to create more space to accommodate the bottom half of the leaf.

☞ This is *not the same thing* as changing the image size, but the two actions are often confused. In Canvas Size you are changing the size of the canvas by adding space around your existing image, or cropping it down.

Go to menu: **Image > Canvas Size...** and in the dialog box double the height number or change the values from inches to percent, and put in 200%. Before hitting **OK** you need to show where this extra space will be added. The graphic box labelled 'Anchor' shows where your source image will be in relation to the new canvas size. Since you want to create space below the existing image, click in the top middle box. The arrows in the other boxes show

where the new canvas will be added. In this case, since we have only changed the height, the additional space will be added below the existing image. Now hit **OK**. Select menu: **File > Save As...** and give your new image a suitable name.

Now you need to collect the bottom half of the page and paste it into the new file containing the top half. Bring the image showing the bottom half of the page to the foreground by clicking on it, or (if you have a lot of images open) by navigating via the menu: **Window > Documents > [filename of file you want]**.

Using your newly-acquired selection skills, select the manuscript page, excluding the dark background. The most important place to ensure your selection is tight to the edge of the leaf is at the top, where it will join on to the upper part. The best tool for making precise selections like this is the Magnetic Lasso, but if you are not particularly anxious to stick the pages together exactly you could use the Marquee or Lasso tools.

Once you have created your selection, copy the content to the Clipboard:

Menu: **Edit > Copy** or keystroke **Ctrl-C** (PC)/**Cmd-C** (Mac)

Go to the large-canvas image you created earlier by clicking on it or navigating via menu, and paste the clipboard contents into it:

Menu: **Edit > Paste** or keystroke **Ctrl-V** (PC)/**Cmd-V** (Mac)

This creates a new layer with the bottom part of the leaf on it, and it sits on top of the background layer, so covers it up until you move it.



Select the 'Pointer' or 'Move' tool from the tool palette (keystroke V). You can now move the newly pasted layer content by click-dragging or, for fine adjustment, by using the cursor keys.

Once you are satisfied with the positioning of the lower portion of the page, you can either save the finished document as a .psd file to retain the individual layers (allowing you to go back later to make further changes), or you can 'flatten' the new layer onto the old one, creating a finished flat image. See below for flattening layers.

Layers

The function that is unique to Photoshop is the ability to create layers on your working image. This is analogous to using successive transparent overlays on an overhead projector. The base image is called 'background' and although you can perform changes and alterations on this base layer, it is better to leave it untouched, and make all your alterations in layers which can be turned on or off (i.e. made visible or not), changed around in order (by dragging) and will be saved when you exit the document, as long as you save in Photoshop format. (TIF format will also save layers, but the file will be very large.)

The use of layers is essential in digital restoration, since actions which cannot be changed when performed on the background can be altered or edited when stored in a layer.



The commonest problems encountered when working with layers all go back to a single error: every time you create a layer, that layer will be the 'active' one, and is selected in the layers palette. Some actions can be added sequentially without changing back to the background layer, but more often than not – particularly when selecting colours or specific areas of the image – you must return to the background layer to do the work: i.e. click to select the background layer before proceeding.

Creating a new layer

You will need to have an image open, as you cannot create layers on top of nothing. You have already created a layer in the previous example simply by pasting data from another file into the file where you are working. However, in restoration you will want to create new layers without cutting and pasting information.

The layers palette default placement is to the right of the Photoshop workspace and near the bottom of the group of palettes. If you cannot see it immediately check the tabs in the palettes that are visible in case it is not in the foreground. Click on 'Layers' to bring it to the foreground. You will see one layer in place already, called 'Background'. This is your base image, and you will see that it has a little padlock icon, indicating that it is locked and therefore cannot be moved.

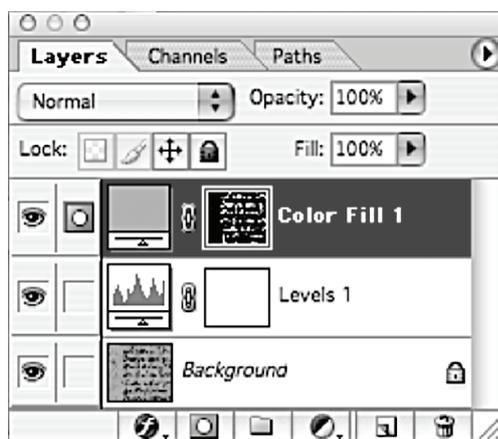


To create a new layer, click on the 'page' icon to the left of the wastebasket. This creates a new blank layer which can be used for drawing (e.g. to annotate the document) or performing an action that you may want to keep separate from the main image.

You can also create a layer by cutting/copying and pasting data either from the original document, or from another (as described above). Using the Text tool [T] also creates a new layer, so that you can edit the text and/or move it around the document after creating it.

To avoid data loss, **SAVE** your document each time you create a layer, and as often as you think of it in between times.

Adjustment layers



In this example two layers have been added to the background: a level adjustment layer, and a colour-fill layer. The eye symbol to the left of each layer description can be clicked to hide or show the layer on the image: if the eye is showing the layer is visible, if not, it is not visible and not acting on other layers..

The blank box to the right of the eye can be clicked to link one layer to another. You can then merge two layers into one, or perhaps move two layers together if they have pasted shapes in them. The icons on the bottom of the window will allow you to create new layers or discard layers. The trash can icon

probably self-explanatory. The icon immediately to its left, like a new page icon, creates a new layer. When you create a new layer it is invisible, but you can draw on it as you would on an overhead projector transparency. To create a new *adjustment* layer – one which has an effect on the layer(s) below it, use the black and white circle (or yin-yang) icon. This is a sub-menu with a selection of types of layer which are discussed below.



Adjustment layers are created by clicking on the bisected circle icon on the bottom of the Layers window. The tiny inverted triangle on the icon indicates that there is a pop-up menu under this icon. Clicking on it will show a list of different types of layer that can be added to an image, mostly with functions that change the appearance of the picture. Selecting one of these functions will bring up a dialog box with a series of options. Those that are relevant to the restoration of images will be discussed in Chapter 3 (Restoration 1).



If you are performing a layer adjustment that affects the whole image, such as a global level adjust, and your image is very large, you can save processing time by creating a low-resolution copy of your master image, doing the level adjustment on the small image, and then dragging the layer from the low-resolution copy to the high-resolution master.

Layer order

Layers only act on those beneath them (i.e. those under them in the list), so layer order can be important. The order of layers can be changed, and doing so changes the way in which the image appears, since layers only operate on those below them. A global level-adjust layer will have a different effect on the image depending upon where in the set of layers it appears. This is one of the reasons it is extremely useful to be able to go back into a layer and change its behaviour or colour values: if you move a layer you may want to edit its contents because it has a different effect on the image, or you may want to edit other layers that are now being affected by it.

The action of a layer on those lying beneath it can be changed from a simple overlay style to something more complex by using the drop-down menu on the top left of the window (which reads 'Normal' in the default view. Use of different overlays is discussed in Chapter 4 (Restoration 2).

Once you have created the desired effect using layers you can either save using Photoshop format, which will retain the layer structure you have created (enabling you to come back and continue editing later). To save as a final version which will be smaller than the .psd version, choose 'Flatten Image' from the Layers drop-down menu located under the  on the top right of the layers window.

This menu also allows you to merge linked layers (linking is done by clicking in the box to the right of the 'eye' box, and shows a chain icon when a layer is linked to another) and 'Merge down', which merges a layer with the layer immediately beneath. In general with adjustment layers this is not used, as the action of the adjustment changes when merged with another layer.

3**RESTORATION 1**

Ethical process
Sequence of operations
Global level adjust
Select colour range
Single-area level adjust
Colour fills
Pattern fills
Combinations
Filters and Variations
Layer transparency
Common errors and problems
Variations on colour select
How far to go
Unethical process

What I'm aiming for here is to create 'quick and dirty' results which give visible (if not very attractive) restorations, but enables you to see the method behind the process and expand on it yourself. When you are working on your own images you will have time to refine colour ranges, tidy up tiny areas that are unsatisfactory and generally spend as much time as you need to get your image looking exactly as you want it.

Here my priority is to restore readability: the result may not be particularly pretty, but if it is transcribable then it has met its purpose.

Ethical process

Generally, when someone sets out to restore an image to readability, they are aiming to retrieve information rather than attempt to re-create the look of the original document. However, there are some restoration specialists (e.g. Fotoscintifica in Bologna) who claim to restore documents to their original appearance, even though this is not actually possible. The work they do is extremely meticulous, and the results are spectacular (if expensive), but they are aiming for something different than the academic, who mainly wants to retrieve lost information. Their work is analogous to the furniture restorer, who makes a damaged antique piece of furniture look as if it was new, as opposed to someone who simply polishes and tidies up an old item so that it is no longer falling apart. The second type of restoration still has its distressed look, but looks aged rather than decrepit.

The sort of restoration undertaken by DIAMM is the second sort: we aim to return a document to usability, but the results are not always beautiful: principally we need to be able to retrieve the information recorded in the text.

At the outset we asked ourselves: where is the line between restoration and just enhancing the view a bit? At what point do we become editors?

The answer is that the first time you touch an image, no matter what the process, you are editing it. If the image is not an adequate representation of the original you may have had to sharpen or colour adjust, but these are editing actions. This is why we insist that the image should be as accurate as possible without any post-process activity.

Doing an 'ethical' restoration means that

- Only colours which are patently not supposed to be there should be removed;
- Only figures about which there is no question can be treated as colours to be brought forward;
- If there is a figure which you believe to be a certain shape you *must not* restore it so that your belief is confirmed. You can only use colours on shapes that are not in question, and you may find that the shape you believed is there is actually something different.

If you intend to publish a facsimile, but have decided that there is not much point in reproducing the source as it is (since it is unreadable), and you will produce a restored version, then what you publish is no longer a facsimile, it is now a photographic edition, and you *must* make it clear in what way your published images differ from the original source. Ideally publications like this should have the original and restored versions printed side-by-side.

Apart from anything else it is important to the library that their source is not misrepresented by your edition.

With damaged sources, publishing in colour gives the reader far more information than grayscale, but if cost restricts you to grayscale you may have no choice but to alter the gray images in a way that you would not have needed to do in the colour ones.

Before you start to restore it is advisable to decide how far you will go in this process. Are you simply going to try to reveal what the machine tells you is there, or are you going to 'help it along' by deciding before you restore what you expect to find? This is very easily done, and poses an ethical dilemma: your restoration may be correct in your eyes, but someone else reading the original document may disagree with the result you have achieved.

The ethical position taken in DIAMM was that any process should be applied to whole areas of a document, not single colour-artefacts that we believed to be musical notation. Attempting to 'brighten up' something that appears like a faded note form is akin to going over it with a pencil in the original source to make it look more like what we think it is. We could be right about the content, but we could just as easily be imagining that it says what we expect it to say. Where we attempt to retrieve a specific colour range we make sure that we select it from point where there is no question that what we are selecting is indeed music.

Basing research materials on digitally restored images lays your work open to criticism, and you need to decide for yourself at a very early stage how much further you will go on from global adjustments into fine detail, which is where disagreement about content will arise.

The best way to present digitally restored materials is in the electronic medium, generally online, where you are not tied to any single view as you are in print. Different layers can be presented as overlays that can be turned on and off for example. The online environment is already offering quite spectacular possibilities for digital editions, with powerful image-delivery, annotation and display facilities becoming commonplace.

Sequence of operations

The order in which you do things may not seem important, and in many cases it isn't particularly important, particularly when layer adjustments can be swapped around, but there is a simple sequence of actions that will 'clean up' an image surprisingly quickly, and may save you a lot of time painstakingly trying to separate out colours that you want to find.

Having created a series of adjustment layers, you may find that swapping the order around gives you different results.

Global level adjust

All the actions in this session will be done in layers as adjustment layers, and the first of these is a global level adjust. All the following images have been pre-cropped and saved to Photoshop format to save time.

This image is extremely dirty as it was used for binding reinforcement and was covered in glue. It has gone rather like a thick piece of greaseproof paper. It is clear that there is music on the page, but it is difficult to read. The colours are all rather similar. The image used is <Sample1.psd>



<p>First of all I 'Save As...' into Photoshop (.psd) format, giving the file a new name. This leaves the original file untouched.</p>	<p>Save As... Photoshop format with new name</p>
<p>Since you are going to be performing processes that will change the colour, and since the scales are saved in the original file, cropping off any data that is extraneous to that which you need to work with minimizes file size and will speed up any actions you apply. The picture above has already been cropped.</p>	<p> Crop (C) to optimize file size and memory usage. SAVE</p>
<p>The first restoration process should always be a global level adjust: this action 'spreads out' the colours in the picture so that the pale colours become paler and the dark colours become darker. Use the adjustment layer tool to add a level adjust layer: click on the bisected circle at the bottom of the layer window and</p>	<p> Global level adjust creating new adjustment layer SAVE</p>

<p>select from the menu 'Levels...'. Moving the white slider to the left makes the whole image paler, which is not terribly useful. However moving the middle (grey) slider to the left makes the mid-tones paler, and the notes can be darkened by moving the black slider to the right. Hit OK to apply.</p>	
---	--

Select colour range

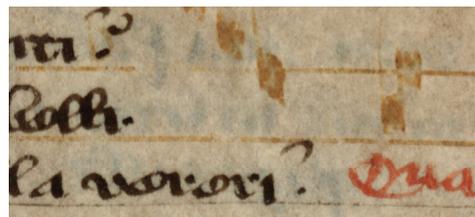
There is a certain trick to colour-selection for restoration that has nothing to do with how clever you are with Photoshop, or with Photoshop's tools and features. For those who have used Photoshop before, the techniques described here are likely to be familiar, and the question may arise: 'but I tried that, why didn't it work for me?'

The key is not in understanding Photoshop (though that helps) it is in understanding the image.

A very high-resolution image gives you a far better understanding of the document than you can obtain with anything of lesser quality, and even by looking at the original document. Although smaller images can be used for restoration, only the high-resolution makes detailed colour selection possible, and it is in accurate colour selection that almost all digital restoration takes place.

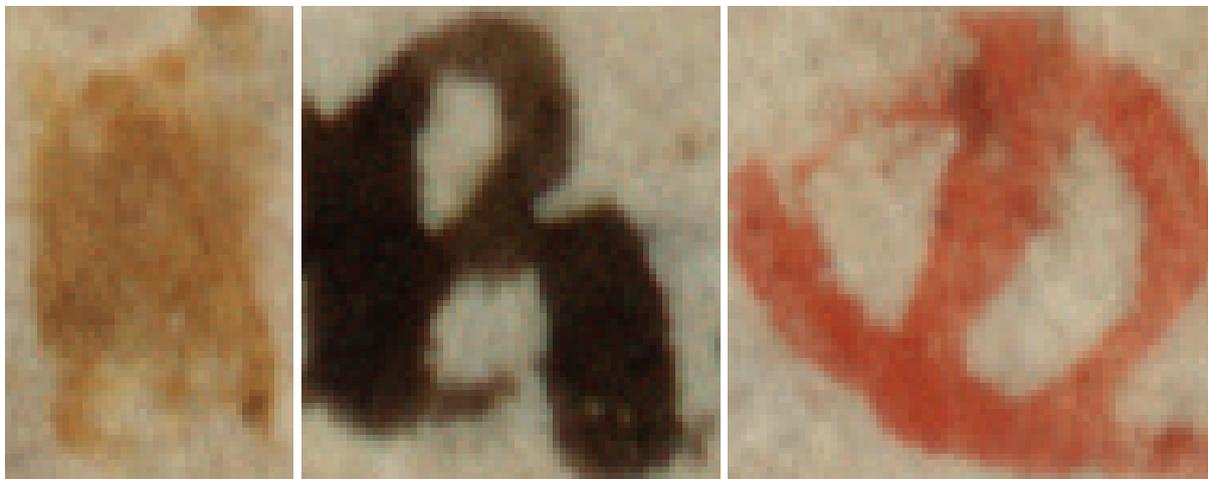
This is a close-up of a section of palimpsest. The image is slightly crooked, but I have *not* straightened it up as that causes a blurring of the fine detail. Superficially we are dealing with four colours:

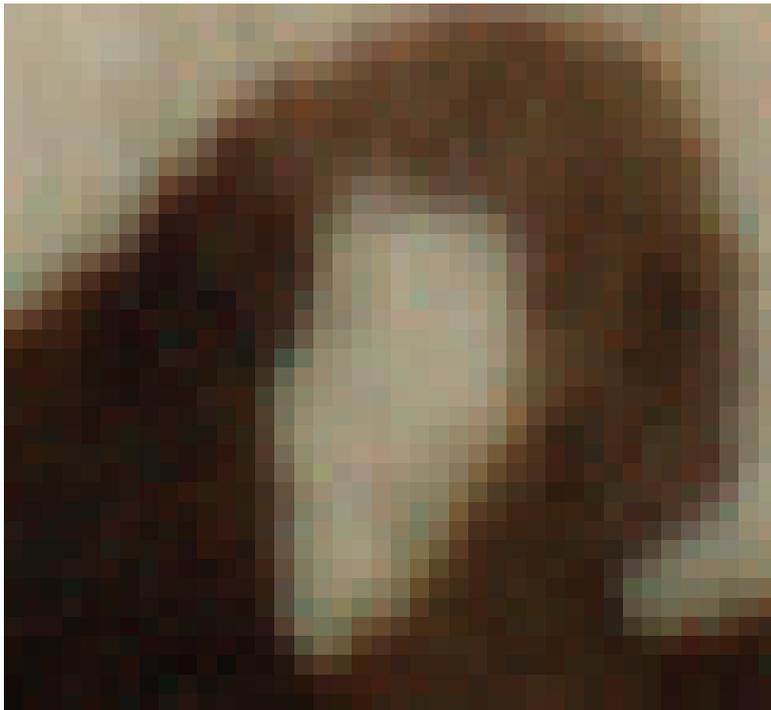
- the background of the parchment;
- the original writing (pale brown);
- the overwriting (dark brown or black);
- a second colour of overwriting (red);
- You may also decide that there is a fourth colour present: the grey of show-through from the reverse of the leaf.



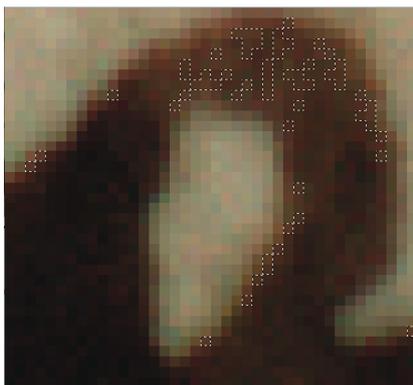
We need to get in a lot closer to see what is really going on: doing so will cause the image to break up into pixels, but that suits our purpose very well, since we can see the difference between the various pixels in the picture.

Because these are digital images, we can zoom up beyond pixel-for-pixel view:





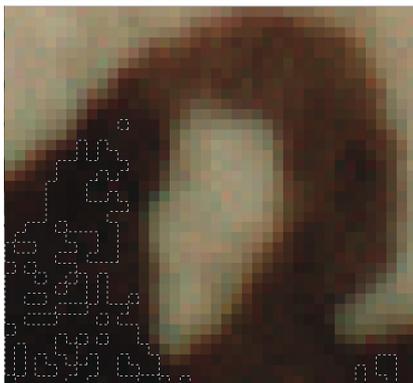
At once you should be able to see that each of these colours is breaking up into many different tones, not least in the background, which is far from being a single flat colour, regardless of any show-through; this is particularly visible with the red and pale brown writing. When you colour select with photoshop you are dealing with electronic information which is far more detailed than your eye can see. Dark brown is not *just* dark brown. In fact even one stroke of the pen has dozens, perhaps hundreds of colours in it. Using the Magic Wand does not show how many colours are available, as its range is not subtle enough, but using Select Color Range is extremely



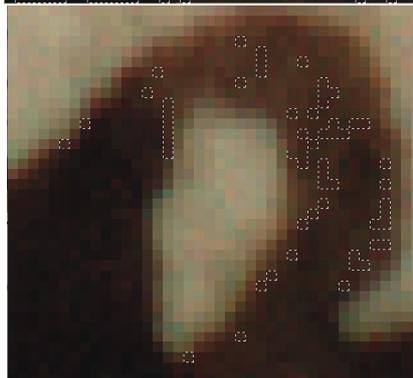
revealing, even with the fuzziness set at 20. This is why, when you select colours, the results are sometimes surprising, and occasionally disappointing. The eye does not see what is actually in the image, hence the need for many colour picks when you are selecting a specific colour, not just one or two.

This is why we never use compressed image-storage formats, as this fine gradation between colours would be lost.

Various colour selections collect only a few pixels, even with a high degree of fuzziness.



Adding a number of selections based on the main part of the letter stroke together, and then replacing those colours with a pattern based on the background shows why, when ink colours are selected and removed, the edges very often remain:



In the next example, a global level-adjust has little useful effect, so I have moved directly on to a colour-range selection to pick up the colour of the music notes and darken them.

The image used is <Sample2.psd>.



<p>The simplest way to select the note colour is to zoom up to 200% view. The image must be magnified to do this, as you need to be able to select specific pixels.</p>	 Zoom to 200% view (Z)
<p>Go to the 'Select Color Range' tool. Set the fuzziness to 20 in the first instance. Use the pointer to select pixels within the note shapes on the image behind the dialog box.</p> <p>It is unlikely that a single colour selection will collect enough pixels for the notes, so keep clicking around on the same note, checking the dialog box preview each time to make sure you haven't selected the wrong colour pixels. You can undo a single selection by using Ctrl/Cmd-Z, but only one.</p> <p>You may get better results by selecting colours in several places on the image. Don't forget also to select the edge colour, as ink bleed is usually paler than the body colour</p>	<p>Menu: Select > Color Range...</p> <p>Set fuzziness to 20, select note colours on image.</p> <p>Pan around image for range of colours, including edge-colours.</p> <p>Ctrl/Cmd-Z to undo a selection that looks wrong.</p>
<p>You will need to balance the note-colours that you select against the fuzziness to get the right colour range. When you think you have an optimum selection, hit OK.</p>	<p>Try altering fuzziness to get a more accurate selection range.</p> <p>hit OK to select.</p>

There will be a short wait while the software finds the colours and creates a selection. This appears as a dotted active line surrounding the selection. At high magnification you can see whether this has been an effective selection or not, but don't give up on it: zoom out before deciding to discard what you have chosen. Assuming the selection is good, we need to darken it up to make these notes more readable:

<p>Create a new level adjustment layer (as above). To make the dark tones in this selection darker, move the leftmost black slider to the right until the image looks how you want it. Clicking in the 'preview' checkbox in this dialog box will show you the image with and without the adjustment you have just done without actually applying it. Hit OK to apply. SAVE your work.</p>	 Create new level adjustment layer to darken note shapes. SAVE YOUR WORK
--	---

Single-area level adjust

In this image the problems are similar to those in Image 1, but I know that a global level adjust won't work on the whole image because of the heavy staining in patches. This staining was almost certainly caused by the use of oak galls in a physical restoration process several years ago. Galls darken up faded ink, but at the time they were in common use conservators didn't realise what it would do to the surface a few years down the line.

The image is <Sample3.psd>; this image is also referred to in Chapter 4.



<p>As in the first image, the first process to try is a global level adjust, adding an adjustment layer. The results on the unstained parts of the page are good.</p>	 Global level adjust creating new adjustment layer SAVE YOUR WORK
<p>The division between the stain and the rest of the page is very clear, so the magnetic lasso is the ideal tool for selecting out the bits of the image that are going to require further work. Once an area is selected any new layer created will only affect the selected portions. Create the selection of the top stain using the magnetic lasso, then add to the selection using <code>shift</code> to select the lower stain.</p>	 Create a shape selection using the magnetic lasso tool.
<p>Once again go to the level adjust tool. This time, because we've already performed a level adjust on the image, the histogram shows gaps in the colours, but it has values right across the spectrum now. A second level adjust may sort this out...</p>	 Level adjust on specific colours

If that is not enough, then the next step is to look at the stained areas more closely and try selecting specific colours to brighten or darken (in this case it may be best to select the background dark colour and make it paler: see variations on colour select, below).

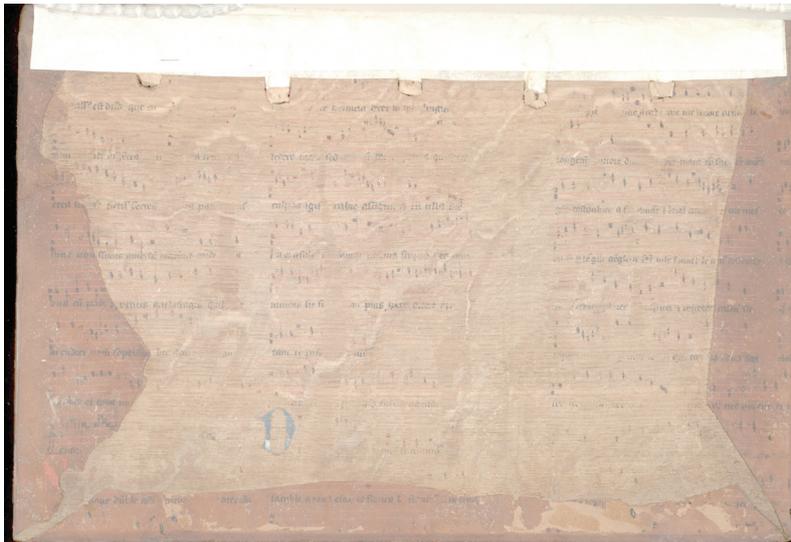
Another example where this type of restoration is very effective is in offsets onto wood. In the example below the offset is all that remains of the original paste-down, which was discarded during rebinding. The whole image is very dark, and there is some difficulty in distinguishing between wood grain and ink offset.

The image used is <Sample4.psd>



Because this is an offset it is a mirror image, so the first action is to flip it to the right way around. This is in the same menu as rotating, but is a different function. Menu: Image > Rotate Canvas > Flip Canvas Horizontal.

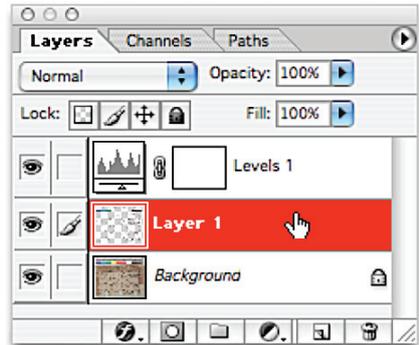
A global level adjust is reasonably helpful, but the point of doing it on this image is to fade back the dark background colour as much as possible. To avoid making the result a garish yellow colour use the white slider a small amount, but make a greater adjustment using the gray slider. Don't worry about making the music notation pale as well as the background for the moment.



Because the music notation was already dark, it's unnecessary to darken it further, however we don't want it as pale as the background is now. Turn off the level adjustment layer you have just made by clicking on the icon of the eye beside that layer. The eye disappears and the level adjust is no longer visible



Turn off the level adjustment layer visibility

<p>Now zoom in to see the music notation clearly. Use the 'Select Color Range' tool to select the black of the ink. Remember to balance between fuzziness and number of individual colours selected. Once you are happy with your selection range, hit OK.</p>	<p>Select > Color Range... to select black ink of notes.</p>
<p>Instead of creating a level adjustment layer, this selection will be used to create a new layer altogether. Copy the contents of your selection, the Paste immediately. Doing this automatically creates a new layer containing the contents of your selection range. It is placed exactly on top of the selection area, so is not obvious yet. On the Layer list you will see a new layer between the background and the level adjustment layer.</p>	<p>menu: Edit > Copy or Ctrl/Cmd-C menu: Edit > Paste or Ctrl/Cmd-V</p>
<p>At the moment, when you return the level adjustment layer to readability by clicking where the Eye icon should be it will affect everything beneath it the same. Move the new layer that you created until it is above the adjustment layer. Click-and-hold anywhere in the layer title box, and drag the layer to the new position (your cursor turns to a hand). Immediately the notes are visible as they are not affected by the level adjustment (which is now below), so they remain dark against the paler background of the main image.</p>	

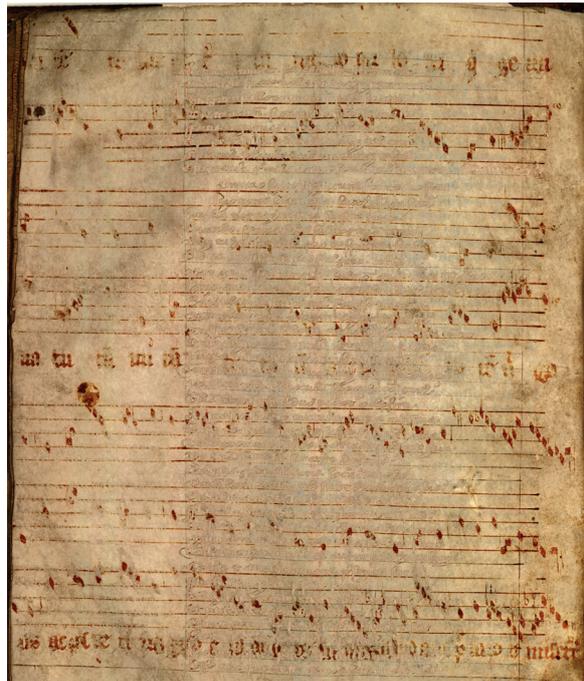
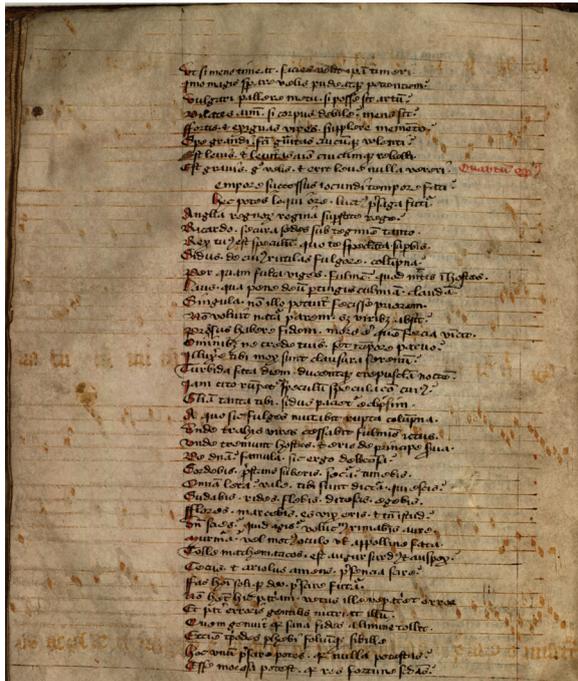
Repeat the process for the red stave lines and blue initials separately, and the result is a much more readable record of the inks left behind on the board.

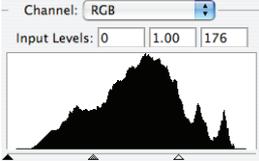


Colour fills

This technique is mostly done when removing dirt or palimpsest over-writing.

This is what I would call an easy restoration sample because the two texts are visibly different colours. This technique is the same that would be applied on text deletions, but the closer in time between the original writing and the deletion, the more difficult it becomes to separate the colours. The image is <Sample5.psd>.



<p>The first action is always to see if a simple level adjust will clarify what is on the page.</p>	 Add level adjustment layer
<p>The over-writing is quite varied in colour, particularly the edges of the words where the ink has bled, and selecting it will require quite careful balance between fuzziness and number of colours selected. It may be necessary to create two or more select colour ranges and treat them separately. Enlarging the view higher than 200% will assist in the process of selecting very specific colours.</p>	<p>Select > Color Range... to select black ink of text. Menu:</p>
<p>Once you have a selection first of all try applying a level-adjust layer just to that selection. This is done just the same way as applying a global level adjust, but as you have pixels selected it will only act on the selection.</p>	 Add level adjustment layer
<p>When the white slider is moved to the left, to make the notes paler, they tend to go bright yellow. This is because the colour selected (in this case) is too dark. The level adjust alteration will not give you the desired result. You can abandon this adjustment and go back to your selection simply by clicking on cancel (keystroke Esc). The dialog box disappears, and you see your selection again.</p>	 <p>Level Adjust histogram</p>
<p>Instead try selecting 'Solid Color...' from the adjustment layer menu. You can choose a colour in the colour dialog box by clicking in the main part of the box, the colour gradient panel to the right, or by entering specific numbers in the RGB or CMYK value boxes. However the most effective way of finding the right colour is to click on the background colour in the image. The colour selection box will change with each click, and you can see whether you have the 'right' colour to create a masking colour fill because the effect is previewed in the display.</p> <p>One important aspect of what you will see by clicking on the background is that it is made up from a huge number of</p>	 Add solid colour layer; click on the image to find an appropriate colour <p>SAVE</p>

<p>component colours. We are used to assuming that the background is a single colour, but even with the eye you can see that across the leaf the overall colour ranges very widely. You may need to click-select on the background many times before you find a colour that works – the best places to try are in the middle of the leaf, between the lines of text that you want to disappear. When you find a colour you like, click OK.</p>	
<p>If you are happy with the result that is good. However, if you think that it has left a ghost of the original text behind, you may need to expand and feather the original selection. Use the History to go back to your 'Select Color Range' selection – you can click back and forth to find the place you want.</p>	 <p>History</p>
<p>Now go to menu: Select > Modify > Expand... and type '2' into the pixels box and click OK. Your selection will expand noticeably.</p>	<p>Select > Modify > Expand... 2</p>
<p>To feather the expansion you just created, go to menu: Select > Feather... and type '2' also into the pixel box and click OK.</p>	<p>Select > Feather... 2</p>
<p>Now go through the same steps as above to create a colour fill layer.</p>	 <p>Create colour fill layer</p>

☞ If you accidentally lose your selection remember that you can get it back by stepping backward in your History.



Pattern fills

The image is <Sample6.psd>. Colour fills are a quick and easy way to deal with something that you want to hide or merge into the background. It works best with small shapes, like the tight secretary hand on the previous sample.

However there are some images where the thing you want to hide is very large – perhaps a large stain or, as in the case of the next sample, overwriting that is very large and obtrusive. The best way of hiding large areas is to select the colour and then do a level adjust to fade the colour back. However, particularly with dark colours, this is often not satisfactory, as the colour gets brighter rather than lighter, and with colours that cover a

wide spectrum a single level adjust on the whole thing will not work. In an example like the one below, a pattern fill is almost certainly going to be more effective, although the background colour is perhaps too variable for an ideal solution.

In this case a single colour fill could work, but the result of having such a large area of solid plain colour may be almost as distracting as the original text was.

It would probably be more effective (if slower) to create a pattern fill based on the background of the document: so instead of finding a single colour among all those that go to make up the background, you are able to take a much broader selection and create a more generalized fill pattern.

To create a pattern fill you need to duplicate the background layer, as this is where the pattern will be created. Having done that, go to menu: **Filter > Pattern Maker...** This shows you a preview of your background image. Find a clear area of parchment without any writing or obvious marks on it. Preferably from the centre of the document, or close to the text that you want to obscure (in this case the large title hand, as this area will be best for hiding the over-writing. Point and drag diagonally to create an approximately square selection. It does not have to be very large. Click on **Generate**. Then at the bottom left of the tools on the right hand side of this window, click on the (tiny) disc icon to save the selection as a pattern. Click **OK**.

Filter > Pattern Maker...;

Select area;

Click on **Generate**

Click disc icon to save as a pattern

SAVE



You can create a pattern before or after you have created your colour range selection. You will not lose the selection by duplicating the background layer and using the 'Pattern Maker' tool.

When that window disappears you will find that your new layer (a repeating pattern) is obscuring everything else. Just turn it off, or discard it. Now you create a colour pick selection as you would normally (as above). Once you have your selection (you may have to create several different selections as the large text covers a broad spectrum of colours), go to the adjustment layer tool and select **Pattern...** from the top group of options. On the patterns palette that pops up the pattern you just created will automatically be selected.

Select > Color Range... to create a colour selection;



Select **Pattern...**

from adjustment layer menu

Combinations.

On an image like <Sample7.psd> you may have to create a number of selections and pattern-fill each one individually in order to get the desired result. This image is graded 'Medium' not because it is particularly difficult, but because it requires patience and meticulous colour selection.

This image has some interesting problems, mainly due to the accretion of different colours of dirt across the surface, coupled with some fading, and the fact that the music is both darker (in the middle) *and* lighter (at bottom right) than the background colour of the parchment.

A global level-adjust is not very helpful here as it creaes a pink cast and does not really clean up the dirt: this is because the range of colours in the dirt and in the notation is so broad, so any blanket adjustment will improve some parts of the document but disimprove others.

In this case the first process is going to involve a colour-range selection, and the most obvious one to go for is the red ink. If you zoom in to 100% (click on the 'Actual Pixels' button on the tool menu-bar) on the cluster of red notes near the top of the page this will give you some idea of what you are going to have to do to make this document readable:



The red notes are straightforward. However the brown notes and the stave lines going horizontally across the page are more interesting. Parchment is a fascinating medium because the ink tends to sit on the surface rather than soaking in. Obviously it does soak into the skin to some extent or we would not be able to retrieve palimpsests, but the majority of it remains stuck to the surface because it is slightly greasy. This behaviour is of course what made it easy for writing to be scraped off and the surface re-used.



In this small sample you can see that the notes range in colour from dark brown, where the ink has survived best, through dark to pale yellowish where it has worn off, down through pinkish to very pale pink where the ink has been rubbed off altogether. The very pale

effect is where the ink has protected the parchment from dirt, but has finally chipped away.

This image is also a good example of poor focus. At 100% magnification you can see that the edges of the colours are slightly soft. A feature of parchment photography is that it is possible to get an extremely sharply focussed picture (assuming the parchment is flat) because the edges of the ink do not bleed as they do on paper. A perfectly focussed image of writing on paper may look soft because of the ink bleed, even though it is possible to see the paper fibres clearly.

Although this image isn't sharp, the colours are sufficiently defined that you should be able to create a useful colour selection.

Using the 'Select Color Range' tool, select all the reds on the page. Remember to zoom in on the picture before choosing 'Select Color Range'. The red is quite easily selected (the smaller the white dots on the preview, the better your selection range), but make sure that you have sampled ink from right across the document: the ink at the bottom of the page is much pinker than that at the top, which is darker and more intense. Use the selection preview options to make sure you have chosen all the colours that appear in the notes.

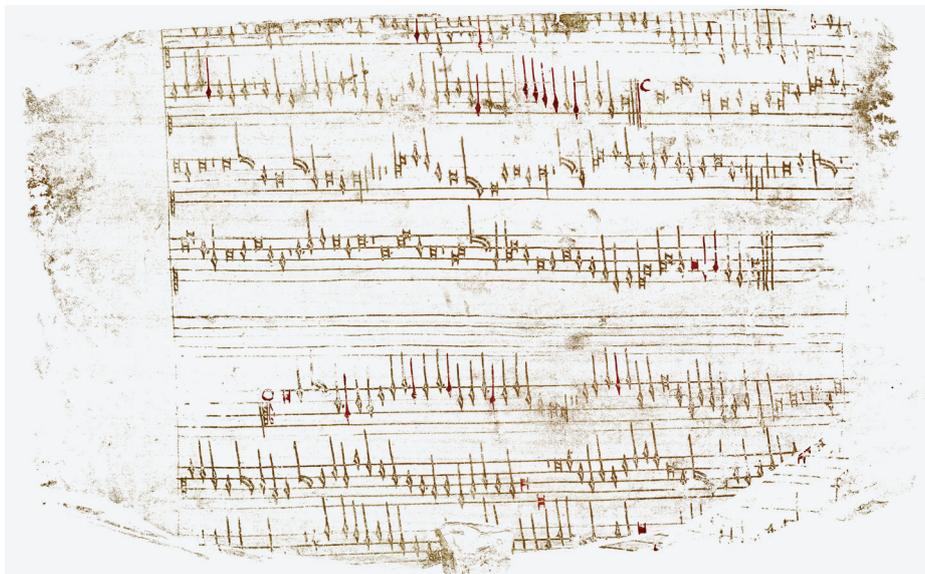


Zoom in to high magnification;

Select > Color Range...: reds across whole page

<p>Since this colour is not faded, there is no need to darken it up. What we want to do is save it so that any other colour adjustments do not affect it. Once you are happy with your selection create a new layer with the red ink only using copy and paste. The new layer appears in the layers menu. You can turn off the background to see the results of your cut-and-paste.</p>	<p>Copy and paste reds, creating a new layer SAVE</p>
<p>The note colours, because of their range, may require several different selections. Because of the overall colour of the document (pinkish rather than grey or yellow) the yellow tones in the notes are probably going to be the most productive colour to select. In the Color Range dialog remember that a tiny portion of white in the preview is probably a good selection, whereas a great deal of white may result in too much material selected.</p> <p>Because the tones across the document are quite similar, keep the fuzziness low (around 10-12), and go for many colour picks.</p>	<p>Select > Color Range...: brown notes (concentrating on yellows); Keep fuzziness low and create a large number of colour picks</p>
<p>As a 'safety' measure I'm going to create a new layer with this selection in it by cutting and pasting. You can name your layers by double-clicking on the layer title name. Now I want to return to the previous selection and try expanding or smoothing it to see if it gives me a clearer selection of the notes.</p>	<p>Copy and paste this selection to create a new layer; SAVE</p>

If the background is made invisible so that only these two pasted layers show, you can see what has been saved. It's quite possible that this type of view would be acceptable as a restoration for some users, since it achieves the goal of readability.



<p>To get back to a colour select that you have used on a colour range previously, without stepping back in the history (in which case you would lose all the interim actions), click on the background layer and then Ctrl/Cmd-click on the layer where you created a selection. The previous selection will reappear.</p>	<p>Return to previous selection: select background, Ctrl/Cmd-click on selection layer</p>
<p>Expanding the selection by 1 pixel picks up a lot of extra detail in the notes that were not caught on the original selection. However when I zoom out to check the whole image it seems also to have picked up a good deal of the edge dirt, which is not desirable. I</p>	<p> Zoom out to check effects of expanding selection range;</p>

can keep it, and go back afterwards to blank out the dirt perhaps with a pattern fill, but it may be better to go back to my original selection and try a different expansion, or just work with the original selection.	Use history to step back if this is not satisfactory
Using History to step back to the original selection, try Smoothing instead. Although this doesn't pick up some missing note colours, it is a better adjustment as it does not find the dirt at the edges.	Try smoothing instead;
A level adjustment layer on this selection doesn't make the notes very much darker and there are specific patches where the pinky note residue has not been found by the previous colour selection, and I will have to select those notes in a later selection. For now though, this is an improvement.	 Apply a level adjustment layer to darken the notes
I'm going to go back to this note colour selection as it seems a useful one: re-select it using the <code>Ctrl/Cmd-click</code> facility. What I'm interested in now is what was NOT selected by this previous selection – in other words the dirty background.	Re-select note colour selection

Because the background is made up from many colours including black and all the way through to pale pinkish-brown, selecting it would probably take many different colour-select actions. By inverting the selection I made to find the brown notes, I will select everything that is not music, i.e. the whole background.

To make this inverted selection use keystroke <code>Cmd-shift-I</code> or menu: Select > Inverse . This will also select the red notes, which were not included in the brown-note selection, but that doesn't matter, as we have saved the original colours of these notes in the first layer created, and putting this layer above all the other adjustment layers will ensure the red notes appear in their unadjusted form.	Invert the note-colour selection using <code>Cmd-shift-I</code> or menu: Select > Inverse
Applying a level adjustment layer seems to make the document go much pinker (probably a deficiency in the colour settings of the original capture), so rather than adjusting the colours individually in the level-adjust drop-down menu, which is difficult and requires quite a lot of tweaking, in this case I'm going to use the 'curves' adjustment layer (on the same menu as levels), which allows a much subtler colour adjustment. This will be described more in Chapter 4. Curves is a different way of looking at colour levels from the histogram on the level adjust. Here the colours are 'described' as a diagonal line, and there is no indication where the colour is concentrated as in the histogram. The top right corresponds to the white slider, the bottom left to the black slider. There is no corresponding point for the grey slider, but you can create adjustment anchor points by clicking on the line; all these points can be moved. I will create an adjustment point in the middle of the line and drag it just a small amount towards the top left corner, which makes the mid-tones in the picture paler. Working with curves or level adjust creates a more pleasing result visually, as colour or pattern fills don't work well across a surface that is so varied in base colour.	 Apply a curves adjustment layer to fade out the background colours Click to create adjustment point on the line, and drag towards upper left corner slightly
All that remains is to colour-select the notes that were not picked up in the first note colour select and darken them. Creating a level	Select > Color Range... for pale

adjust layer will do this, but in this case a more effective solution is to cut and paste the note colours selected in this process (creating a new layer), and changing the layer action from 'Normal' to 'Multiply'. Layer actions will be discussed further in Chapter 4.

notes; create new layer using cut-and-paste; Move layer above curves adjustment layer; change action to 'multiply'

The result is a great improvement, and makes the text perfectly readable. I could further refine my restoration by colour-selecting more of the dark background dirt and fading it out, but actually what has been done already is really sufficient to create a readable result.



If you like what an action has done on part of your image, but not on the rest, you can select and delete the parts of the layer you don't like *as long as you have the layer in question selected*, but only on certain types of adjustment layer (those created by prior colour selection or other selection range)

Filters and Variations

There is a tendency in image restoration to get bogged down in details. If you feel you are in danger of losing sight of your goal, or becoming frustrated by the results you are getting, try something completely off the wall and see what you get. The results can be very surprising, and often very effective.

Filters cannot be added to the image in layers like adjustment layers, so in order to apply a filter that you can turn on and off you first need to duplicate the background layer. The filter is applied to the duplicated layer. Filters are found in the 'Filter' menu.



Unfortunately once a filter is applied to a layer it cannot be edited or the settings changed. The only way to do this would be to create another duplicate of the background layer and re-apply the filter with different settings.

An unsharp mask heightens edge contrast, which can brighten up faded letters on a dirty page. The degree of mask should be checked both with the image at high and low magnification, as what looks good in one view may look terrible in another. A very small

unsharp mask will just add punch to a picture when printed, as it compensates for ink bleed. However over-sharpening makes a picture become grainy as shown in Chapter 1.

Also, using a low-resolution copy of your image to save processing time, try the following filters on duplicates of the background layer, for unusual effects that often heighten readability in unexpected ways:

Filter > Stylize > emboss creates a 3-D effect based on colour intensity

Filter > Stylize > find edges

Filter > Other > minimum ('fattens' all the dark-coloured pixels on the image)

Filter > Render > Lighting Effects (to create a raking light effect)

Filter > Other > High pass

Layer transparency

A layer does not need to affect those layers beneath it at 'full power'. A layer can be made more or less opaque using the transparency option at the top right of the layer window. The number in the box can be changed from 100% by typing a new number into the box, or by using the slider which is accessed from the drop-down menu icon just beside the percentage box.

Turning the picture into a negative of itself with an adjustment layer ('invert') is often extremely revealing, as dark data becomes light and vice versa. The change in what this does to your eye and what it sees often makes previously unreadable material readable, and can be used in conjunction with other layers or adjustments to get an optimum result. Mostly though, this is used to rest the eye and brain and look for information that might be enhanced in less dramatic ways for other users to see.



Another way of dealing with material with very little contrast is to create a duplicate of the background layer, change the overlay effect to 'Difference' and then shift the top image by one pixel left or right and one pixel up or down using the Move Tool (keystroke **V**) and the cursor keys. The filter effect will give a black image when exactly overlaid on the original, but when shifted creates a highlight where there is a difference between the colours in the two layers, effectively finding the edges of all the dark/light components of the image. Again, this is probably not a result for the finished restoration, but could give you valuable information.

If an adjustment layer looks too dramatic to be of any use, try changing the opacity to fade it back before discarding it altogether.

Overlay modes are dealt with in more detail in Chapter 4

Common errors and problems

The commonest error in attempting to create colour selections is to start the selection process without having the background layer selected. You may want the selection result to appear higher up the list of layers, but you can always move and adjust it after you have created it.

Poor colour selections which take up too many or too few pixels are usually the result of a) incorrect balance between the number of colours selected and the fuzziness (or Tolerance if using the Magic Wand Tool).

The more you work with Photoshop images, the more your actions fill up the memory buffer. If you go too far without saving you will receive an error message when you do try to save, telling you that there is not enough memory to complete the operation. All the actions between your last save and this point will be lost. **SAVE SAVE SAVE.**

Losing the mouse

Not as strange as it sounds! With the tools set to precise cursors the mouse pointer can disappear all too easily on a surface that is complex or predominantly pale colours. A quick way to make it more visible is to hold down the **Ctrl/Cmd** key and the cursor will change to a pointer enabling you to locate it.

Variations on colour select

In the example of a very dirty document above I dealt with the background by selecting the note colours, and then inverting that selection to create a selection based on the full background area.

However, if the notes themselves had been more difficult to pick out, it is sometimes easier to create a colour selection based on the colours that you do not want (e.g. the background colours) and then either fade these colours back, or invert the selection and darken what remains. Fading the background is one of the more effective techniques for water stained documents.



Simply making the text you want to read darker creates an improvement immediately, but you can create a more dramatic, and often more readable result by making the same colour select, but replacing the colours you have selected with another, completely different, colour. This is because darkening text nevertheless keeps it in the same colour spectrum as everything else on the page. Replacing the text colour with red, blue, mauve or green lifts it away from its background quite differently, and is often far more striking and thus readable. In the example below (which was originally palimpsest) the two images have had exactly the

same colour range selected, but in the second I have used a replacement colour rather than a level-adjust darkening, as in the first picture.

How far to go

Any painter will tell you that a large portion of the skill they exercise is in knowing when to stop. If you are working with a large batch of images you may be forced to stop. Once you can read it, that's far enough. However, some actions will improve one part of the document, but disimprove other parts, and you may want to create a sequence of restorations, not just one, and use them all side by side.

Digital restoration is analogous to furniture restoration: you can over-restore quite easily. It is not necessary to remove all traces of dirt, nor is it necessary to make faded ink black: only do as much as you need to do to read the text.

Remember that digital restoration is an editorial practice, and must be described as such in any presentation of restored images, not matter how 'accurate' the results are believed to be.

Unethical Process

This may seem a strange heading, but there are processes possible in Photoshop which create superb results on the most severely damaged documents but, in my opinion, are unethical as restoration procedures. It might be more appropriate to describe this as 'Editorial Process; rather than 'Unethical Process', but it is important when using this tool to keep in mind all the time that you are falsifying data.

One of the most difficult types of damage to restore is burn-through of ink, most often found in paper sources from the late 15th and early 16th centuries. In the example below some parts of the text are very clearly show-through and can be eliminated with a clear conscience. However, other damage, where the front and reverse writing coincide, requires a decision about what note-shape has been lost so that a copy can be superimposed.

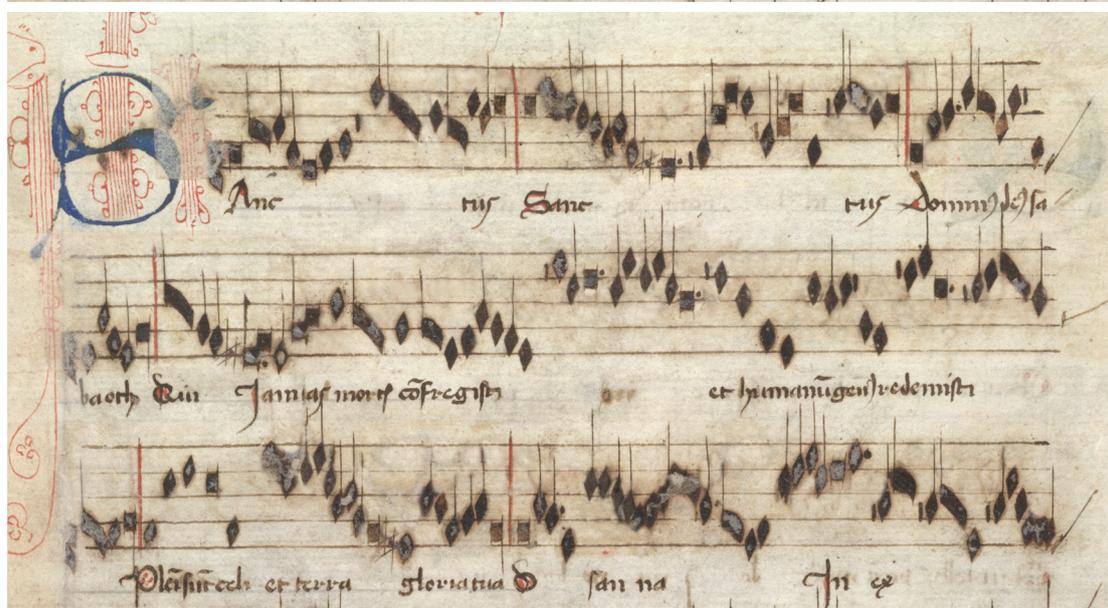


The tool used in this process is the Clone Stamp Tool on the Tools palette (keystroke S). This tool allows you to instantly cut and paste sections of a document onto itself without going through the process of selecting, feathering, cutting and pasting.

You click on a point from where you want to clone, which is your anchor point, by **Opt-click/Alt-click**. Then release the key and wherever you click and drag on the page, the material from the anchor point will be copied.

On the Tools Options Bar you will see a drop-down menu where you can specify what size of brush to use and whether it has hard or soft edges, an airbrush option, and whether the cloning should read from all layers or not. If you create your cloned material in a new layer it can be switched on and off.

In the example below it was generally clear where text was show-through, either from the colour or because of the duct of the hand (slant), and it was simple to delete this where it did not coincide with the surface text. In some places I knew what should be there, but had to clone it from less damaged notes, and there are definitely some possibilities for being wrong in these areas. Where I was really not sure what I was seeing, I simply did nothing, leaving the final decision to the editor of the music, who can work much more easily on these few awkward corners if all the obvious problems have been eliminated in advance.



Once you have a cleaned-up version of one side of the leaf, it would be possible to flip it and superimpose on the reverse as an aid to eliminating the burn-through on the other side. However that would require that both images are the same size and at exactly the same rotation. Even if this were not the case, the tools available in photoshop would enable you to correct one image to match the other.

4

RESTORATION 2

Using levels to assess images

- White and black points

More on Layers

- Layer modes
- Curves tool
- Sharing layers between images

Channels

- The different types of information
- Viewing Channels
- Using Channels

Masks

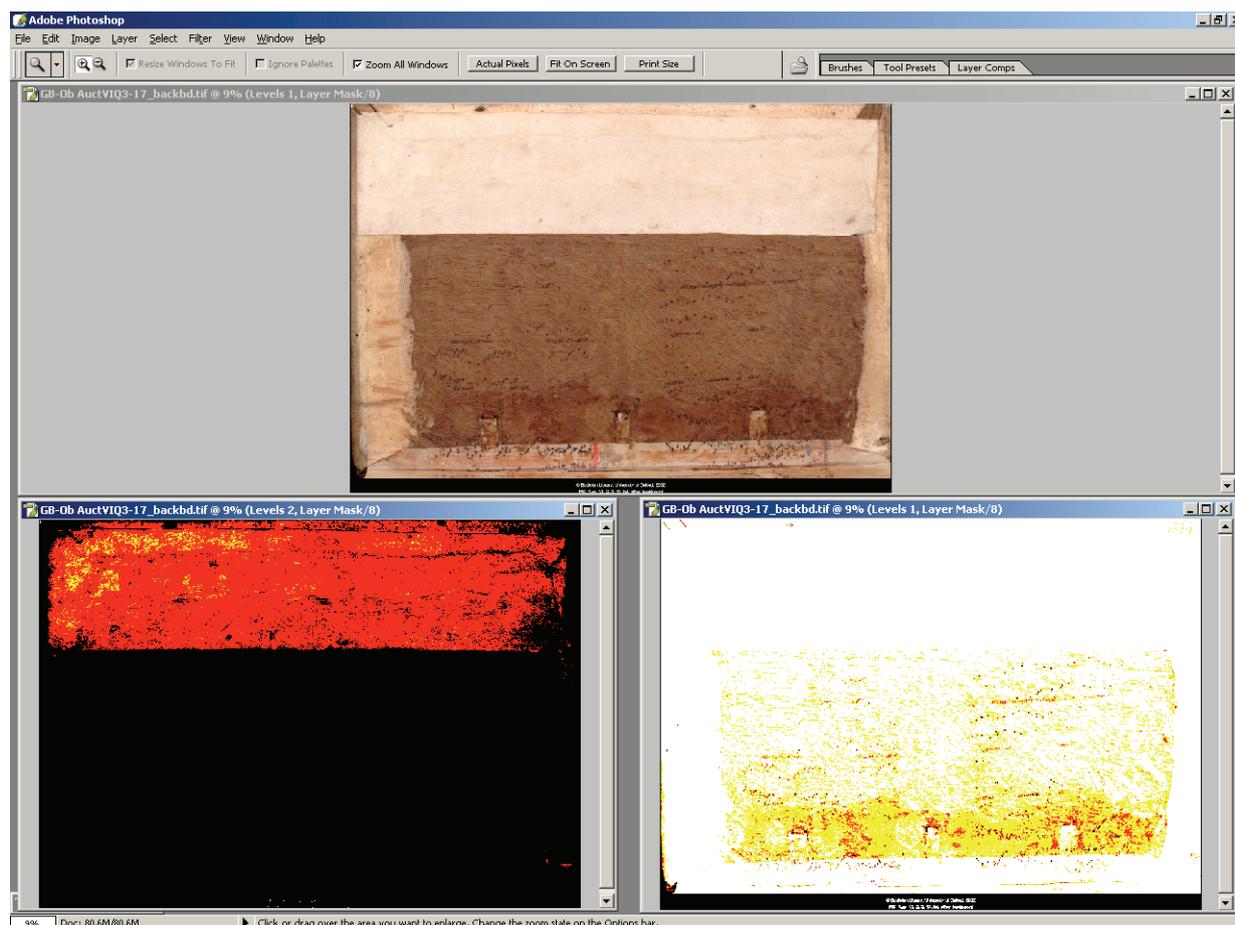
Unsharp masking

Using Levels to assess an image

It's not uncommon to come across images that have been lit unevenly and this can make what should otherwise be a relatively levels adjustment somewhat tricky. Remember we're often dealing information which is just discernible from the background. It can also be difficult sometimes to spot exactly where the brightest and darkest parts of such an image are. Create a Levels adjustment layer just above the background, or a duplicate thereof, (yes, you don't have to use an adjustment layer but just in case you click OK...) and, in the Levels dialog box, hold down the **Alt** key while dragging the white level slider to the left, this will show the lightest pixels first. Note that the pixels will show in their colour i.e. if the red channel is the lightest then red pixels will show before green or blue. To find the darkest pixels in an image drag the black slider while holding down the **Alt** key, this time the darkest colours pixels will show first. This technique shows up uneven illumination fairly well. A similar effect can be had with a Threshold adjustment layer but it ignores the colour information. If need be we can then go on and create a mask/transparency layer based on the luminance of the image (see Channels later).

The image on the next page shows the use of the **Alt** key with the levels sliders to show the lightest and darkest areas of the image.

On the workstation see images <Levels_light_dark.png>, <RGB_difference_1.png> and <RGB_difference_1a.png>.



Remember that the middle eye dropper in the Levels dialog will adjust whatever colour is clicked on to a neutral grey, so if you have an image that you suspect has a not-too-obvious colour problem you can run a quick check by clicking on the grey on the colour card you shot with the subject (you did, didn't you?) and observing what happens to the overall colour balance. **Cancel** out of the dialog so as not to save any changes, unless you want to of course.

More on Layers

Layer modes

One of the advantages of using layers is that the way in which the layer interacts with the image (and other layers) below it can be altered. The two main ways we are most likely to use in the restoration process are the blending mode and opacity.

The blending modes fall into a number of grouped functions, again we are most likely to use just two of these (though others may well have their part to play, particularly the Contrast group). These two groups are the darkening and lightening modes.

The Darkening modes consist of:

- Darken: lighter pixels are replaced with darker ones when the two images (i.e. current layer and those below) are compared.

- Multiply: darkens the entire image, useful to reinforce faint highlight detail in overexposed images.
- Color Burn: increases contrast and darkens image.
- Linear Burn: combination of Multiply and Color Burn. Dark pixels can be forced to pure black.

The Lightening group consists of:

- Lighten: darker pixels are replaced with lighter pixels in the comparison of the two images.
- Screen: lightens entire image, used to bring out shadow detail in underexposed images.
- Color Dodge: black values are preserved while there is a decrease in contrast in pixels lighter than 50% grey.
- Linear Dodge: combination of Screen and Color Dodge, light areas can be forced to pure white.

As mentioned in Restoration 1 the strength of the effect of the active layer can be adjusted by the Opacity slider in the Layers palette.



Note that the effects of layers is cumulative, if the effect of one adjustment layer isn't strong enough then you can duplicate that layer to increase the effect, repeat as necessary. If the last duplication is just a touch too strong then use the opacity slider to fine tune the effect. In this way we can get 1.5x, 1.7x, 2.3x or whatever factor we require by a combination of compounding layers and adjusting their opacity.

Open the image <LayerModeMultiply.psd>. As you can see parts of this image are a little overexposed: while we can't restore information that isn't there we can make the best of what we have using another layer with blending mode set to Multiply. There are other 'issues' with this image that we may be able to deal with if time permits.

Open the file <LayerModeScreen.psd>. You will notice that there are two colour samplers placed on the image (yes, Photoshop very usefully saves the samplers in its PSD files). The first sampler is in the upper left corner of the page, the second is on the upper half of the white paper. Make sure both adjustment layers are turned off (they should be already, this is something else that Photoshop saves the state of). Now make the 'Levels adjusted' layer visible, try and assess the image and make a note of the two sampler values. Turn this layer off again. Repeat the exercise with the 'Screen mode' layer. Do you notice any difference?



A 'gotcha': when using some tools (e.g. Magic wand, Paint bucket, Blur and Sharpen) on images with layers a little check box appears on the options bar for 'Use all layers'. What this does is pretty much what it says but it can all too easily lead to a tool not working as you expect, common problems are:

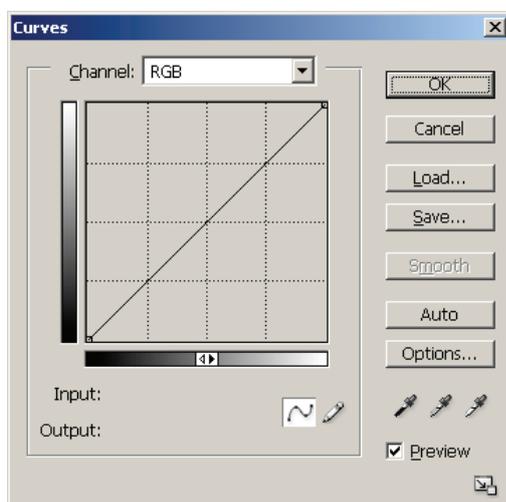
- you have just created a new layer and you click with the Magic wand to make a selection only to have the whole image selected. The tool is working correctly it just that the 'Use all layers' box is unchecked – it only takes the sample from the current layer. To correct this just check the box.
- you want to blur background detail, you have already isolated the foreground objects on a separate layer, you have even remembered to select the background layer (copy, as it's a destructive process) but for some odd reason the tool is blurring your foreground objects too – in this case the 'Use all layers' box is checked when you don't want it to be.

Once you get used to the power of layers it is all too easy to build up quite a collection in an image. This is fine, if you have the RAM!, but can lead to a confusion of layers (“Now just what does this one do?”). There are a couple of things we can do to help reduce this confusion. One is something you will have seen this in most of the example images and that is layer naming – the other is layer sets.

Layers are given a default name by Photoshop e.g. Layer 1; Levels 4; Curves 3 etc. to rename the layer to something more descriptive you can: right-click on the layer (PC) or double click on the layer name (Mac) or go to the palette options button  (top right of the palette well) and choose Layer Properties, this will bring up a simple dialog box where you can edit the name.

When you have more than a few layers it can be useful to organise them into groups related by function: this is where ‘layer sets’ comes in. To use this feature create a new set by clicking on the ‘Create Layer set’  button at the bottom of the layers palette, the new set will have a default name (starting at ‘Set 1’) which you may wish to change, the process mirrors that of renaming layers. You can now create new layers within that set by selecting the set before creating the layer. Existing layers may be added to the set simply by dragging and dropping (note that this appears not to work with an empty set for some, no doubt logical, reason). Quite often this leads to the dropped layer not ending up in the right position in the set but, just like ordinary layers, they may be dragged to a new position. Sets can also be created in, or dragged to, another set. A very useful feature of sets is that the whole group of layers within it can be disabled by clicking on the ‘eye’  icon to the left of the set. Oh, and to tidy up the palette sets can be ‘collapsed’ so that only the set name is showing by clicking on the triangle icon to the left of the set name, clicking it again opens the set up.

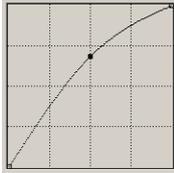
Curves tool



This tool, like Levels, is used to re-map pixel values to new ones. The main difference is that where the Levels dialog has three sliders (plus input and output levels) the Curves tool can have up to nine sliding points. This does make it a more complex tool to get to grips with, the bonus for doing this is the flexibility it gives you. Unfortunately the dialog box is really too small for such a tool but this gives us an opportunity to look at how keyboard control can make things much easier. But first off we’ll take a quick look at the dialog box and *some* of its settings.

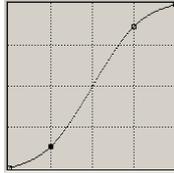
The image here is the Curves dialog as it first opens. Note that the ‘curve’ is a straight line, at the moment there will be no effect on the image (you can easily see effects if you ensure that the ‘Preview’ box is checked). Most of the controls are similar to the ones you are familiar with in the levels control. The Options button leads on to a sub-menu where tool defaults can be set (and reset) and other options such as snapping neutral midtones, clipping levels and so on. The Auto button uses some of these settings to make a ‘best guess’ adjustment to the image. It may give you what you after but it is much more likely that you’ll need to take manual control.

In this image there are no numbers by the Input and Output labels as there is no control point for them to echo the values of (excepting the points at the end of the curve).



Clicking anywhere in the central grid places a control point on it that pulls the line to it. In the example on the left the midtone values are made lighter, the effect here is like moving the midtone slider in the levels control to the left – so far no major advantage.

So let's take a look at how this tool might offer benefits over levels...



Here a second control point has been added and the first one moved up and to the right. Now we have a curve! This curve will have the effect of: lightening and compressing highlights, darkening and compressing shadows while enhancing midtone contrast. Since we can have up to nine control points the affected the potential for both finely-tuned correction and strange effects is high (try moving a control point significantly out of position). These points can be on a per-channel basis not just RGB.

Note that the active control is denoted by being solid black instead of an outline.

This tool allows you to bring out detail in narrow sections of the tonal range in a way that levels can't manage, though it does make it trickier to use (think multiple gamma controls). Part of its trickiness comes from the control points not being very easy to control(!) with a mouse so we'll have a look at how the keyboard can make this easier...

To start with there is accurate placement of the initial point, if you're not quite sure where in the tonal range to place it you can click on the eyedropper tool (it's in the options bar for this tool) and click somewhere in the image that contains a tone you want, you'll notice that a control point appears on the curve at the tonal point you have just sampled, release the click and it disappears – not necessarily what you want. To fix the point on the curve press **Ctrl/Cmd** while you click.

Once you have the point(s) you need then moving it with the mouse gives only fairly coarse control, to make things easier click on the point you wish to move, now use the **arrow** keys on your keyboard to move the point in the direction you want to move it. Note that the numbers by Input and Output change by one for each press, using the Shift in conjunction accelerates the change to 15.



These number-change and 'accelerator' keystrokes work in the same way in any number box, though the accelerator is usually 10. It is often a much easier and more reproducible way to change values than using the mouse, another advantage is that it's easier to keep your eye on the image and not get led astray by histograms etc.

The three eyedroppers on the bottom right of the dialog box set black, grey and white points respectively. The grey point eyedropper can be very useful in neutralising colour casts: click on the Set Grey Point eyedropper to activate and then click on a point in the image that contains a known neutral grey (preferably middle tonal range), depending on how well you chose the reference grey the colour cast should be removed.

There is a lot more to the Curves tool but hopefully, especially if you've not felt comfortable with it before, this short introduction will encourage you to experiment with it.

Sharing Layers between images

It quite often happens that when there are problems with an image, other images shot at the same time have the same (or at least similar) problems. If this is the case then try and pick the most representative of the group and work on getting that right. Once you have the adjustment layer(s) providing the corrections you're after these can be easily shared with the other images – saving a lot of duplication of work.

To share the layer(s) have both the image containing the layers to share (source) and the destination image open, make sure that both are visible and, with the source active, then quite simply drag the required layer(s) from the Layer palette onto the destination *image*. That's it. If the effect isn't quite right it may be that adjusting the layer opacity (and/or duplicating the layer) will give sufficient fine tuning.

The same process can be applied to layer sets too.

 Note that adjustment layers are scaleable, they don't care about differences in the aspect ratio or size of the two images. This lends itself to speeding up work on big images: downsample (a copy of!) the image and work on this, note that this will be for overall adjustments not fine detail. Because the memory requirements can be drastically reduced the processing can run somewhat faster, especially if there are lots of layers. When you've finished, save the image for safekeeping, and just copy the layer(s) across to the full size version as above.

Channels

To quote Adobe: "Channels are greyscale images that store different types of information."

What are these different types of information and how can we make use of them?

Starting with almost the simplest image: the single channel just contains the greyscale data for that image. If we move on to basic RGB colour images then the 4 channels 'created' contain the values of each of the colour channels, 3 of them show the greyscale version of either Red or Green or Blue data, the fourth channel shows the combined (or composite) data and is shown in colour.

Note that for CMYK images an extra channel is needed making a minimum of five.

Another type of channel is the alpha channel, these allow you to create, manipulate, save and load masks. These are used for protecting areas from, or enabling areas for, edits. They also are used to create transparency effects by allowing pixels to show through or not according to the pixel values in the alpha channel. Note that alpha channels are greyscales, this means that the 'strength' of a particular effect is applied according to these greyscale values (black means no effect, white means full effect).

Viewing channels:

 You will need to have an image open in order to do this.

Open the channels palette by either clicking on the Channels tab in the layer palette if open or **Window > Channels**, if you prefer to have both the Layers and Channels visible at the same time, then you can click-drag the Channels palette out of the Layers palette. Assuming a simple RGB image you can click on the eye icon to the left of the channel (this works in the same way as the eye icon in the layers palette) to view that channel's information – note they are greyscales. You can also click on the R, G and B channels in turn (after first clicking on RGB if its not already active), to view another channel in combination with the first, say R and B, then hold down **shift** while clicking on the second channel. To view all channels in combination just click on the RGB channel. As usual there are useful keyboard shortcuts to change channels: **Ctrl/Cmd-~** shows all channels, **Ctrl/Cmd-1** shows Red, **Ctrl/Cmd-2** shows Green and **Ctrl/Cmd-3** shows the Blue

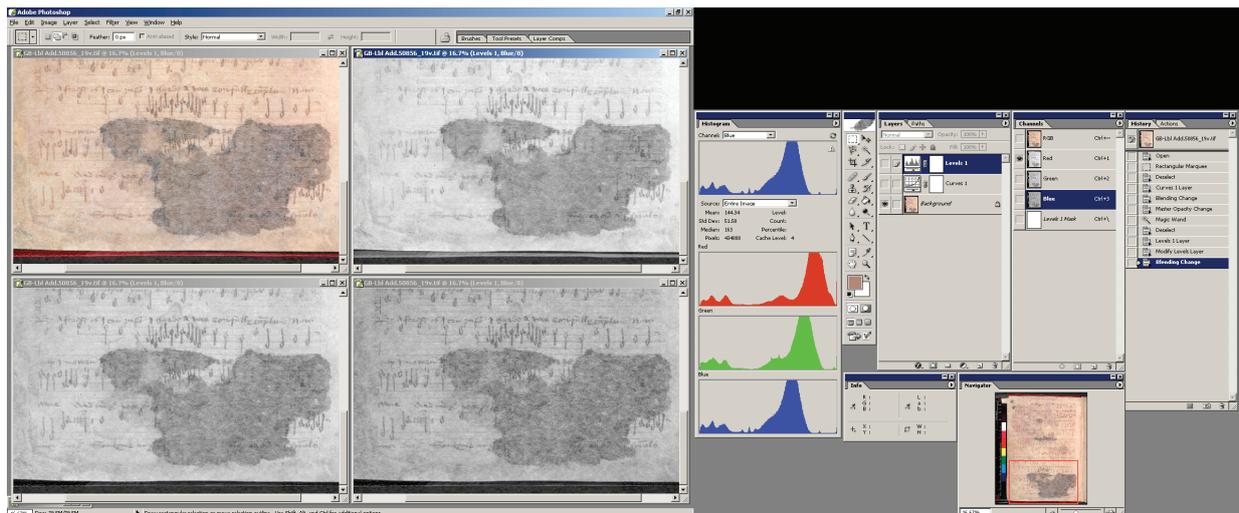
channel. Using the keyboard in this way makes it much easier to observe any differences on the image between channels, as your eyes can more easily stay on the subject.

☞ The **New Window** for <image name> option on the **Window** menu allows you have more than one view of an image, these may be: two (or more, this applies to all the following) different zooms; two areas of the image; two individual or combined channels and so on.

These views may be synchronised for zoom (not relevant in first example!) and/or location. Go to **Window > Arrange** and choose which option you want.

How can we use Channels?

Sometimes image information can be seen more readily in single channel view than in the RGB in combination. For example if a dark stain is predominantly a mixture of red and green then the blue channel may show objects (or show them more clearly) that are obscured by that stain in RGB (see the example on the next page: compare the RGB and B channels in the stained area). It may then be easier to use this channel for making selections than the combined view. Also more advanced manipulations (easier in later versions of Photoshop) can be used to extract the data from one channel and use it to 'rebuild' data in other channel(s) – a sort of cross-channel cloning.

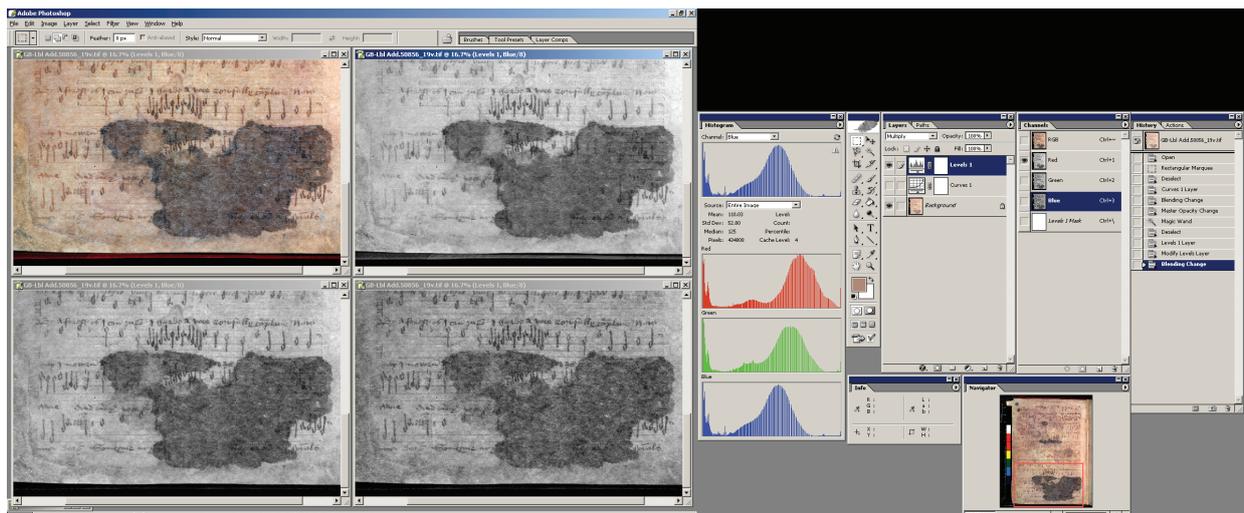


RGB_Difference_1

This is the image as it was shot – the adjustment layers are turned off.

Note the overall shape and position of the histograms (this is Photoshop CS so the R, G, B are viewable separately).

The dual screen setup allows the tools and palettes to be placed on the second screen allowing more room for the image itself. The dark area at top right is due the second monitor running at a lower resolution.



RGB_Difference_1a

Now a levels adjustment layer has been turned on: the black point has been moved up in order to darken the image.

Look at the histograms again, note the gaps that have appeared due to tonal stretching.

The four views of the image were created by menus:

Window > Arrange > New Window for <image>, then

Window > Arrange > Tile, clicking on the appropriate channel and finally

Window > Arrange > Match Zoom and Location.

A number of filters are also quite happy to operate on single or double channels. For instance, consumer digital cameras are particularly prone to noise (graininess) in the blue channel and one of the ways of ameliorating this is to blur the blue data – this is easy to do, select the blue channel (and thereby disable the red and green) and **Filter > Blur**, or for more control **Filter > Gaussian Blur**.



The blue channel rarely contains *detail* that we can use as the eye does not respond well to these wavelengths (and not at all in the very centre of the eye). This is made use of in lossy image compression algorithms such as JPEG, PhotoCD etc. where the blue information can be discarded to a much greater extent than either the red or green data

Masks

Masks, as their name implies, protect (or enable) areas of an image from (or for) editing and viewing. As they are greyscale images they can be used to alter the strength of the effect of other tools: black means no effect, white means full and mid-grey somewhere in between. They can be created in a number of ways:

- Using **Quickmask**  (found on the main Tools palette) after making a selection with any selection tool
- Creating a mask from a Threshold adjustment layer
- Painting directly onto an alpha channel (click the **New Channel**  button in the Channels palette)

Selections were covered earlier in this workbook; couple them with channels and the power and versatility increase enormously. Selections can be made into masks, masks into selections. To save a selection: click on the **Save Selection**  button in the Channels palette. Conversely to load a mask (alpha channel) as selection: click on the channel you wish to load then click on the **Load Selection**  button.

As has already been mentioned these are greyscale images, that means they can be operated on by any tool that operates on greyscale images: sharpen; blur; add or remove noise; curves; levels and so on.

Painting on masks is an easy way to refine selections, again there are a number of approaches to this. The easiest is to use the Quickmask tool:

- Create a selection using an appropriate selection tool
- Click on Quickmask  (press **Q**), a rubyolith overlay will appear (its colour and transparency can be changed).
- Paint with any of the painting tools to either add to the mask (black), subtract from the mask (white) or a mid-grey to alter transparency.
- To return to Standard mode editing click on the **Standard mode**  button (or press **Q**). Note how the mask has become a selection again with the edits incorporated.

Having refined your mask you may want to save it as a Channel Mask so that it is not lost and is available for further editing.



It is a matter of personal preference but for a number of people graphics tablets make certain aspects of image work very much easier. They are much more 'natural' to use for painting, drawing, outlining (as with the Magnetic Lasso) and other such tasks than a mouse. In addition, most are pressure sensitive and can therefore alter some aspect of the tool's behaviour as the user changes pressure – with the brush tool, for example, the most common is to increase the width, increase the saturation of a colour or a combination of the two.

Masks can interact with another mask/selection. For example we may have may a selection that has very distinct borders (e.g. the edge of a page) that we want to mask from the same set of operations as another area that needs the the boundary between it and the rest of the page softened. In this case we would create both selections separately with the appropriate tools, save them as masks and then combine them.

Adding these two masks together load one as a selection (click **Load Selection** button in a channels palette), hold down **Ctrl+Shift** (PC) or **Cmd+Shift** (Mac) and click the channel you wish to add. Of course this new selection can be saved as a mask too (there is a limit of 56 channels). Mask combination is not restricted to addition, they can be subtracted and intersected as with the selection tools themselves.

Layer Masks are masks associated with a layer and are used to restrict the effects of an adjustment layer to areas where the mask is transparent.



Note that a Layer Mask may not show in the Channels palette unless either the mask already existed as a Channel or the layer with which the mask is associated is active.

There is a lot more to channels, masks and layer masks than this workbook attempts to cover. For some idea of their power and flexibility open an image containing an adjustment layer, make a selection (anything will do as you won't be saving the image) then **right-click** (PC) or **alt-click** (Mac) on, firstly, the Layer Mask Thumbnail and look at the

range of options; secondly **right-click** (PC)/**alt-Click** (Mac) on the Layer Thumbnail and note the range of options here. Close the image or save under another name and start exploring the range of options available.

 A Layer Mask can be temporarily disabled by **shift-clicking** (PC/Mac) on its icon in the Layer palette. **Shift-clicking** again will re-enable it.

As mentioned earlier in the section on layers we can usefully create a mask based on the luminance of an image, this can be used as a way of dealing with uneven illumination problems. To create such a mask click on the RGB Channel in the Channels palette or press **Ctrl/Cmd+~** then hold down the **Ctrl/Cmd** key and click the RGB channel in the Channels palette. You now have a luminance-based selection which can be saved as a luminance mask, used with an adjustment layer (or both.) Note that this mask can be manipulated like any mask to enhance, diminish or otherwise vary its effects.

Unsharp Masking

This subject has intentionally been left until last – it is the final operation you should do to an image (other than save the file). There are two main reasons for this: firstly, you won't know how much sharpening you will need until you know the output size and destination (printer or screen); secondly, and more importantly, the process is destructive.

Photoshop has a few sharpening tools, they're found under the Filters menu (**Filters > Sharpen**), they consist of

- Sharpen,
- Sharpen Edges,
- Sharpen More,
- Unsharp Mask.

We'll ignore the first three as they are blunt instruments with no user settable parameters. The Unsharp Mask however does have settings that the user can vary (as indicated by the ellipses after the menu entry). The settings are:

- Amount: the extent to which the image will be sharpened, too high a value leads to halo effects around edges.
- Radius: the 'area' that will be sampled to calculate the sharpening effect. Setting this too high will also make the sharpening obvious.
- Threshold: the difference in value that adjacent pixels need to be before they are seen as an edge. Use higher values on noisy images so that the noise does not get emphasised by the sharpening process.

The Unsharp Mask only works on a single layer at a time so you may need to merge layers (or flatten the image) to achieve the effect you're after – of course it may be that you only want to sharpen selective bits of the image. You can also do this by creating an appropriate selection and using this to define the area to which the effect of the filter should be limited. It is also possible to apply the Unsharp Mask to selected channels, so we could sharpen just the red and green channels while leaving the blue channel as it is (which, as you may remember, tends to be noisier).

A couple of 'general rules':

- images for print will need more sharpening than those destined for screen use. The only real test here is to make proof prints on the printer the image is to be printed on – remember to turn off print driver enhancements as these can mask the changes you are trying to assess.
- the more an image has been reduced in size the more sharpening it will need to give the *appearance* of detail.

you need to view the image at 100% to see the effects, any lower zoom factor will almost certainly lead to the image being oversharpened.

The easiest way to learn about the interaction of the settings is to adopt the process of changing just one parameter at a time, pushing it to extremes (use the Preview option if your computer can update fast enough), then change parameter and repeat.