

Using 'raw' files from digital cameras



Conrad Taylor is Chair of the Electronic Publishing Specialist Group of the British Computer Society and teaches Advanced Photoshop and other courses in electronic publishing.

This briefing paper has been written in response to growing interest among digital photographers in using the so-called 'RAW' files that can be downloaded from some higher-end models of digital camera. **Conrad Taylor** describes what they are, what advantages they have over JPEG files, and what you need to be able to process and use them.

Why JPEG isn't good enough

Most users of digital cameras are content to have the camera store their files in the widely-used JPEG format; consumer cameras offer no other choice. The main advantage of the JPEG format is that the images are compressed, so you can store more images on a memory card or on your hard disk. JPEG is also a well-documented standard, and the files can be opened and dealt with by a very large range of applications.

The best-known disadvantage of JPEG is that the form of compression it uses is *destructive* – the term often used is 'lossy' – and once the damage is done, there is no returning to the quality of the uncompressed image.¹

What is less often realised is that digital cameras perform other kinds of processing on captured image data. These processes are based on engineering assumptions built into its onboard software, and may not suit what the photographer had in mind for the image.

'RAW' = unprocessed

Many professional photographers want a greater degree of control over how the captured image is processed, analogous to the darkroom manipulations of developing and printing a film negative, in which the range of contrast can be manipulated intelligently and for full creative effect. This is provided for by the increasing number of digital cameras which offer the option of saving 'RAW' files, which have hardly been processed in the camera at all, and have certainly not suffered destructive compression.

Essentially, the idea is that you transfer the 'RAW' files to your computer and process them afterwards, with the benefit of more accurate and flexible controls and better visual feedback in desktop software, and then save them in whatever file format you like.

Because of the widespread use of the capitalised form 'RAW', you might be led to think that RAW is a standardised format, like TIFF or JPEG. In fact, each manufacturer of digital cameras has its own proprietary format for the storage of unprocessed camera data: see table in the left margin. For the rest of this paper, I'll refer to these files as 'raw' (lower case), in the sense of being unprocessed.

Some raw file formats of camera manufacturers

| | |
|----------|-------------|
| Canon: | CRW and CR2 |
| Fuji: | RAF |
| Kodak: | KDC |
| Minolta: | MRW |
| Nikon: | NEF |
| Olympus: | ORF |
| Pentax: | PTX and PEF |
| Sigma: | X3F |
| Sony: | ARW |

1. For details, see my PowerBriefing on [JPEG Compression](#).

Understanding image processing

To understand what it is that you give up when you allow the camera to save a processed file, it is worth examining what is involved in a camera's internal workflow. We may divide this into a series of steps:

- Image capture and digitization
- Amplification of the captured values
- 'Demosaicing': Interpolation of camera sensor data into an array of RGB pixels
- Making adjustments to the values of those pixels (brightness, tonal curve, colour balance, saturation etc)
- Subsampling the pixel data to 8 bits of precision per colour channel
- Application of digital sharpening
- Compression of the image, using JPEG's discrete cosine transform methods
- Writing a JPEG file to memory: the compressed image plus some camera metadata

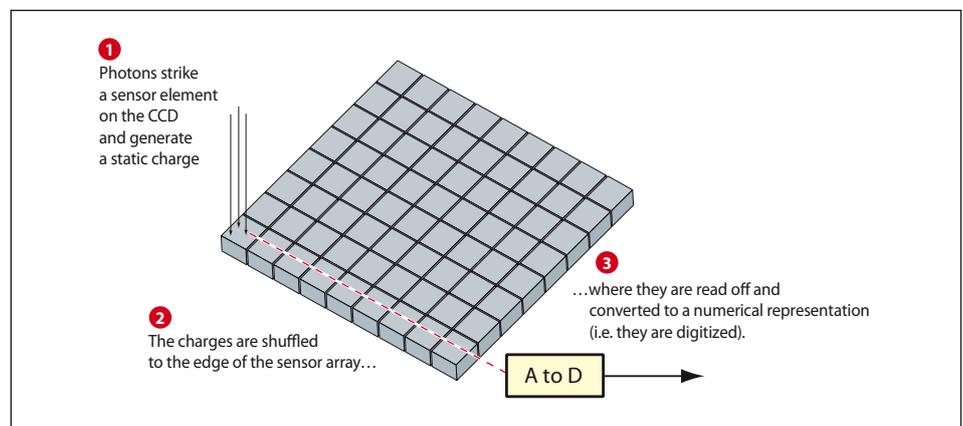
From light to numbers

As light streams into a digital camera lens, it is focused onto a rectangular image sensor, typically a CCD (Charge Coupled Device). Most digital camera chips present an array of square light sensors arranged in rows and columns.²

The CCD array is made up of tiny light-sensitive capacitors. When these are bombarded with photons, static electrical charges build up in them: the greater the number of photons that strike, the higher the resulting charge.

These capacitors are linked together (hence 'coupled'); this provides a means whereby, after an exposure has been taken, the resulting charges can be shuffled to the edge of the imaging chip. Here, the strength of charge recorded by each sensor is read off and turned into a numerical format – the image has now been digitized (see Fig. 1).

Figure 1.
How the CCD converts photons to a static electrical charge, and then to a numerical (i.e. digital) representation.



2. There are other patterns in use, as in the Fuji Finepix cameras which have elements set in a hexagonal array like a honeycomb.

Recording colour

Without some kind of colour filtration, a CCD array would record only the levels of brightness in an image: the result would be greyscale data.

There are four possible approaches to this problem. One solution is to capture three images in succession, each taken through a different colour filter, rotated into place. This method gives superb resolution, but can be used only when photographing still objects. Some specialist cameras used by art galleries for digitising their collections work on this principle, such as those built for the National Gallery in London.

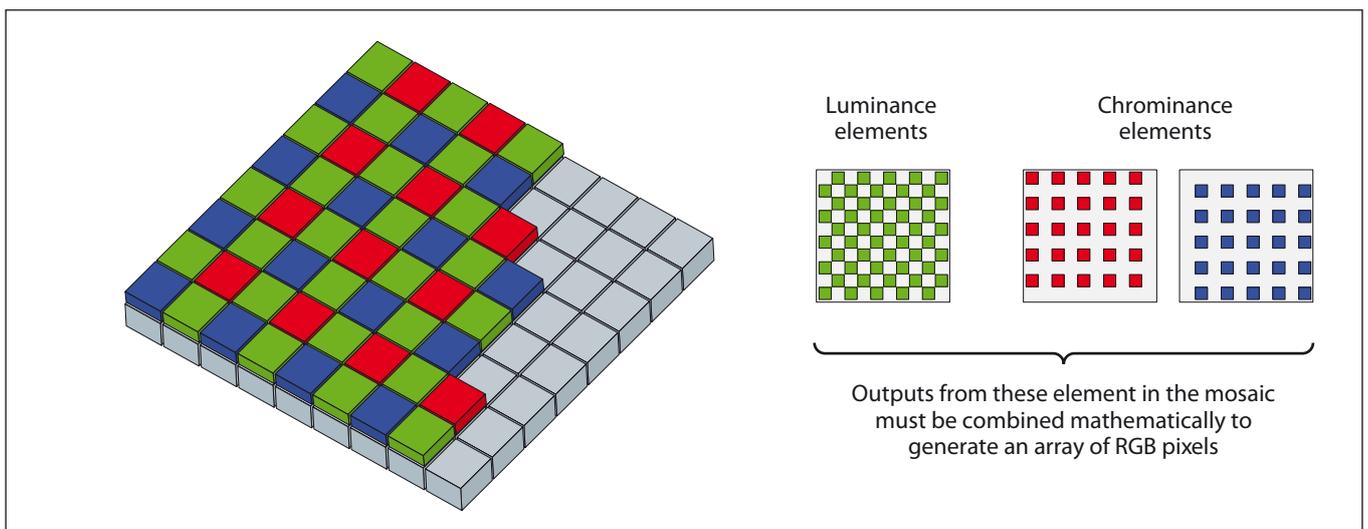
Another solution is to use a precision prism as a 'beam-splitter', directing three copies of what the camera sees to three separate, colour-filtered CCDs. This is the approach used in high-end video cameras. However, a 3-CCD system complete with beam-splitter at the high resolution required for stills digital photography would be very expensive, and also bulky.

There is one innovative colour sensor on the market that has light sensors for red, green and blue stacked on top of each other at each location on the chip, exploiting the fact that these different colour of light penetrate silicon to different depths. The Foveon X3 'direct image sensor' is at present used in a few models of digital camera from Sigma Corporation (and some specialist technical cameras from Hanvision).³

The colour-array approach

Every other kind of digital stills camera uses a single monochrome imaging chip like the one illustrated on the previous page, covered by a colour filter array (CFA). The most common filter array system is called the Bayer Filter (see Figure 2), named for its inventor Dr Bryce E Bayer of Eastman Kodak, and this is the one we'll use for the purpose of explanation.

Figure 2.
Bayer colour filter array, the most common single-chip imaging system used in digital cameras.



1. Foveon Web site: www.foveon.com.

Human eyes are most sensitive to light in the green part of the spectrum, and in the Bayer system, half of the sensors are green-filtered. The levels of brightness which they record are used to define the range of brightnesses in the image: Bryce Bayer called them the *luminance-sensitive* elements. Data retrieved from the red- and blue-filtered elements are used to determine the colours in the image: Bayer called these the *chrominance-sensitive* elements.

We should note in passing that although the Bayer filter pattern is the most common in use in digital cameras, it's not the only game in town. Sony, for example, have brought out cameras using 'RGBE' filters: red, green, blue – and what they term 'emerald', similar to cyan. Their publicity argues that the response to colour of an RGBE-filtered sensor is closer to that of the human vision system.

For our purposes, all we need note is that the pattern of raw data retrieved from an RGBE sensor will be different from that from an RGB Bayer filter sensor, and will have to be interpreted differently.

Taking a real-world camera: Nikon D200

For the purposes of explanation, it helps to take one particular model of digital single-lens reflex camera as an example; I have chosen the Nikon D200 principally because I use one!

Described as having '10.2 effective megapixels', this digital camera uses a Bayer-filter image sensor chip with 3872×2592 elements. The analogue-to-digital conversion process records brightness levels from each image element at a precision of 12 bits – that is, 4,096 levels of brightness are recorded at each element position.

By pressing the 'QUAL' button on the top of the camera and rotating a selector wheel, the photographer can choose to save images at between one of three levels of JPEG compression, or as 'RAW'. Alternatively, two versions can be saved at the same time: a raw image file together with any one of the three JPEG options.

A quick boost – and there's your raw file

After the image has been captured and digitized, the first stage of processing it is to boost the recorded brightness values. The photographer can choose to operate a Nikon D200 at various levels of 'sensitivity', using the analogy of film ISO numbers: from ISO 100 to ISO 1600. Using a higher ISO setting lets you take pictures in dimmer lighting conditions. Of course, choosing a higher ISO setting does not change the physical properties of the image sensor chip; it merely instructs the camera to amplify the recorded values to a greater degree.

Despite great advances in the design of CCD light sensors, they are still inherently somewhat 'noisy' – that is, they produce background rate of random error which pollutes the accuracy of measurements recorded. If a lot of light has fallen on the sensor while an exposure was taken, the signal to noise ratio is very good, but if we compensate for low light availability by



Figure 3. Nikon D200 body with 80–200 mm zoom lens. The buttons which control image storage quality and ISO are to the left of the viewfinder.

Figure 4.

Seven steps to a raw image file.

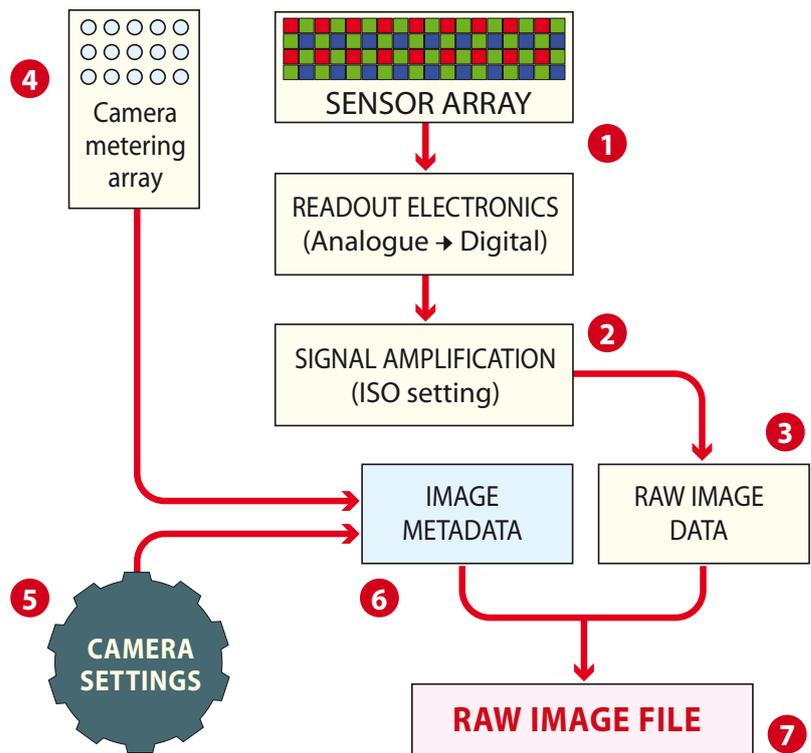
The image captured by the sensor array is converted to a digital representation [1] and boosted in accordance with the chosen ISO setting [2]. This constitutes the raw image data.

In parallel, the camera has measured the scene's lighting characteristics through the metering array [4], and when this is combined with the camera settings [5], this forms an interpretive framework for the image data values which we can call 'image metadata' [6].

Other metadata records such things as the shutter speed and aperture, date and time, lens focal length etc.

The raw image data and the metadata are combined into a file format which we call the 'raw image file' [7].

This may be compressed (albeit in a non-destructive way), and encrypted.



boosting the signal a great deal – say, by instructing the camera to work at ISO 800 – the noise will be amplified together with the weak signal, and the picture will appear 'grainy'.

After signal amplification, if the D200 has been set to save a raw image file, the image data is written into memory as a Nikon NEF file, and no further processing will be done inside the camera. The size of each NEF file is slightly more than 15 megabytes.

The steps so far are represented by the diagram above – Figure 4.

Demosaicing

The next stage in processing the camera image, whether it occurs inside the D200 camera or by processing a raw file on the computer, is to interpret the interlaced Bayer-filter mosaic pattern of 10,036,224 red-, green- and blue-filtered values into an equal number of image pixels, each of which will now have its own complete set of red, green and blue values. The process is known as 'demosaicing'.

Demosaicing is a problematic process, because missing values have to be interpolated from the measured ones, doubling the number of green values and quadrupling the number of red and blue values. Some typical problems that result from figuring out what should stand in for the missing data are: (a) a loss of image sharpness, which is usually compensated for by digitally sharpening the image afterwards and (b) colour fringing artefacts around any sharp-edged objects in the photograph.

More accurate interpolation methods require more complex calculations to be performed, and the designer of a demosaicing system is faced with the

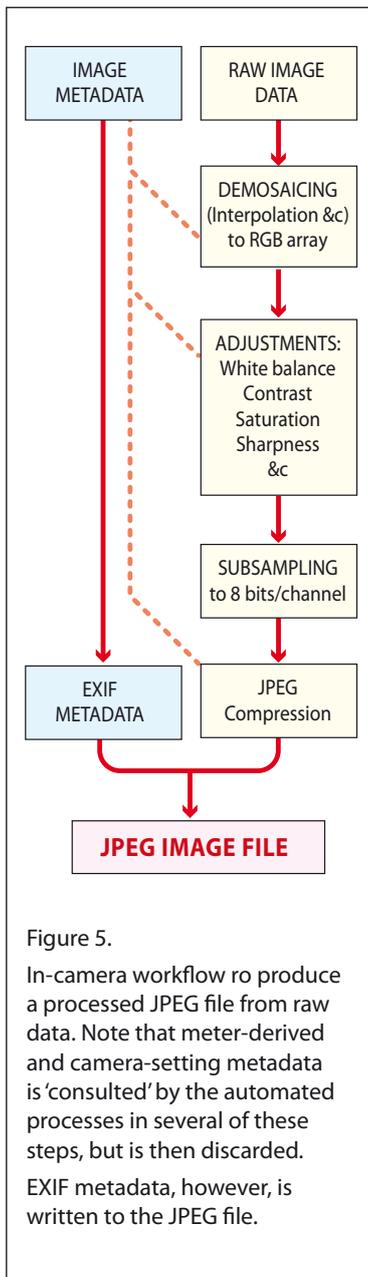


Figure 5.
In-camera workflow to produce a processed JPEG file from raw data. Note that meter-derived and camera-setting metadata is 'consulted' by the automated processes in several of these steps, but is then discarded. EXIF metadata, however, is written to the JPEG file.

problem of balancing the benefits of performance over accuracy – especially if these calculations have to be performed in camera, with the photographer waiting to be able to take the next shot!

One argument in favour of using 'raw' image files is that demosaicing is performed at a later date, on the computer. Potentially, because the need for speedy calculation is not present, the demosaicing algorithms used by the processing software on the computer can be superior to those used in the camera; and, as the science progresses, future processing software could use even better algorithms. (And it is always easier to upgrade software or use a different package than to upgrade the firmware inside the camera.)

Foveon argues that their X3 chip captures a true RGB image, thus avoiding the need for demosaicing and its attendant compromises.

Adjustments and subsampling

As noted above, the image data in a Nikon raw file was gathered at 12-bit precision, but in order to save the photo as a JPEG file, this has to be reduced to 8-bit precision for each colour channels, by a process of subsampling.

Prior to taking this step, the camera will apply a number of adjustments to colour, brightness, contrast and colour saturation, based on measurements made by the camera's metering systems at the time when the photo was taken. (In the case of the Nikon D200, these measurements are collected from a 1,005-pixel RGB colour & exposure metering array.)

The image is also digitally sharpened in order to offset the softening effects of the Bayer interpolation, and other smoothing filters such as low pass filters may be used.

Naturally, these adjustments are also made using a number of inbuilt engineering assumptions – which may not actually match the demands of the actual picture or the photographer's intent. For example, in strong sunshine it may not be possible to records the entire range of light values adequately, and the photographer may want to decide to retain the highlight values at the expense of some shadow detail. If the camera makes the decisions and then throws away 4 useful extra bits of precision, the result may well be worse than if the photographer had had the chance to hand-pick the required adjustments while processing a raw camera image file on the computer.

JPEG compression

The compression scheme devised by the Joint Photographic Experts Group and defined in the ISO standard ISO 10918-1 is the one most frequently used in saving digital images to camera memory. It is a method that involves dividing the entire image into blocks of pixels, typically 8×8 pixels, defining an average value for the block, and then recording how each pixel deviates from the average value by a variable degree of precision.

By reducing the precision with which these deviations are recorded, the image can be compressed a great deal, but at the cost of destroying fidelity to the originally recorded image data. JPEG compression typically records

colour data less accurately than brightness data as a way of saving on file size. All these problems are sidestepped by saving as a raw file.

The D200 camera offers three levels of JPEG 'Image Quality', dubbed Basic, Normal and Fine. A 1-gigabyte memory card can hold approximately 430 images at Basic quality, 220 at Normal quality and 110 at Fine quality.

Image metadata

There is a standard for image metadata, called EXIF, defined by a standards committee representing Japan's digital camera manufacturers. Some of this metadata records when the photo was taken – and even, if the camera can record it, where on the earth's surface! It can also record such things as the shutter speed and aperture and focal length, the white balance and the kind of exposure metering that was used.

Some of the metadata is useful for recording the colour characteristics of the camera. This is important because the numbers which record values within the image are meaningless unless they are pinned down to some kind of reference colour space, such as an InterColor Consortium (ICC) colour profile. The Nikon D200 allows the photographer to choose between two colour spaces: sRGB and Adobe RGB. Having a colour profile definition embedded in the image makes it easier to ensure that a printing device can produce a picture that is faithful to the original (for example when you plug your memory card into a printing booth at a photographic shop).

What's in a raw file?

Each camera manufacturer has its own idea of what should be recorded in a camera raw file, but obviously it must include the data that was recorded from the chip and then signal-amplified. The actual contents will differ according to what kind of physical layout the chip has, what kind of mosaic, etc. This can even differ within one manufacturer's range of cameras: Nikon's NEF has undergone modifications to accommodate different generations of digital camera.

EXIF metadata is also added, as described above. Additional metadata is added in the case of raw files, as shown in Figure 4 above: this records the measurements taken by the camera as the picture was taken, and some computed results of those measurements which the camera would have used in processing the raw image had it been 'asked' to do so. This metadata is useful because when the raw file is opened on the computer, it can be used by the software to provide the photographer with what is usually a sensible starting-point for customised decisions.

Raw camera files may also be compressed – but always in a 'lossless', non-destructive way. Nikon NEF files are not only compressed but also encrypted.

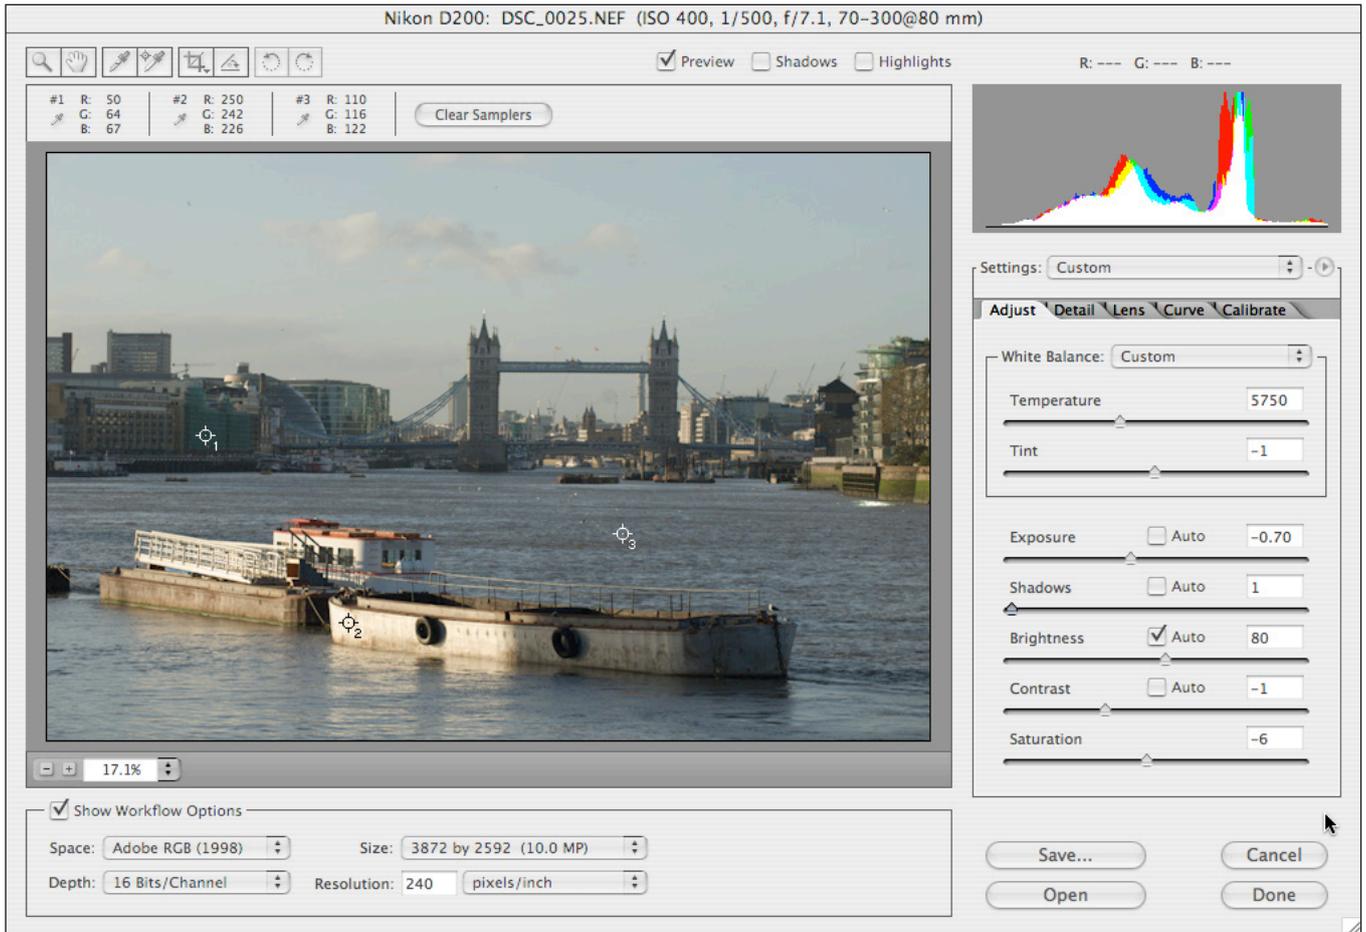


Figure 6.

Opening a Nikon NEF file using the Adobe Camera Raw 3.6 plug-in for Adobe Photoshop CS2.

Note image modification controls on the right, and the histogram display above that.

Also note how three 'samplers' have been placed over the image, and are generating RGB read-outs above the picture as the settings are adjusted.

Using a raw file

The first step in using a raw camera file is to transfer it to the computer, which can be done either with a utility provided by the camera manufacturer, or digital photo library software which recognises the raw file format (for some years I have been using iView MediaPro). The next stage is to open the raw file for 'processing'.

For Photoshop users, this is most conveniently done by adding a plug-in to the 'File Formats' plug-ins folder. Nikon provide a NEF plug-in, but I find that this gives very little control – essentially it is a means of converting from the NEF format to Photoshop 48-bit RGB, using the conversion parameters suggested by the camera.

A great deal more control is by using Adobe's own Camera Raw plug-in. The version of which I am using at the time of writing (version 3.6) supports the raw formats from over 120 different digital cameras, including 32 from Canon, 16 from Olympus, 19 from Nikon, 8 from Kodak etc.

Figure 6 above shows the interface that is provided when opening a Nikon NEF file from a D200 camera using the Adobe Camera Raw 3.6 plug-in. As you can see, the plug-in is a fairly substantial application in its own right, with five panels of settings to adjust to control how the demosaiced file will be adjusted for exposure, brightness, sharpening, colour fringing and so on. When this display first pops onto your screen in response to an attempt to open a raw file in Photoshop, the various sliders are set to initial positions

indicated by the analysis that the camera has performed, and recorded into the raw file's metadata – but one can make adjustments to override the camera's 'thinking'.

Making adjustments requires faith that you aren't damaging the image in any way, and the Camera Raw interface gives three kinds of feedback. Firstly, there is the on-screen display of the image – though you need to know your screen is reasonably calibrated, and remember that the gamut (expressive range) of a colour monitor is not as great as real life, or the range that the camera can record.

To supplement this subjective visual view, there is the histogram display which updates as you make adjustments, and the very useful facility of being able to identify a number of locations in the image for 'Samplers' – in the example shown, I have set one in a shadow zone, one in a highlight and one over the water. The RGB values on a scale from 0 to 255 are also updated as you make changes, and this is a useful way of making sure that you don't burn out highlight detail, for example.

Once you have made your adjustments, you can choose whether to have the raw data interpolated down to an RGB file with 8 bits per colour channel (24-bit colour), or upwards to 16 bits of precision (48-bit colour), and you can then either save the file back to disk as a TIFF or Photoshop file, or simply open the file in Photoshop – in either case, converted in accordance with your preferences.

Specialist tools for raw image processing

Photoshop CS2 is a generalist tool for handling images, which many people already have; Adobe Camera Raw is a great way for pulling raw files into that environment. But this workflow suffers from drawbacks, the main one of which is that you get a single chance to make decisions about how the image will be processed from a raw state to one that Photoshop can natively handle in its workflow.

What many photographers would like is the ability to continue to play with the conversion settings in a 'non-destructive' editing environment, comparing a number of images from a given shoot, and saving not converted images files but rather recipes for how to convert files.

These wishes are becoming true. Nik Software, the soft side of Nikon, has developed an application called Nikon Capture NX which is dedicated to the processing and adjustment of Nikon NEF files.

Both Apple and Adobe also have professional photographers in their sights as potential markets for software, and Apple have come first to market with Apple Aperture, now at version 1.5 – see Figure 7 at the top of the next page.

Meanwhile Adobe Systems is busy developing a similar application called Adobe Lightroom, which at the time of writing is available as a Public Beta until 28 February 2007 so that photographers can offer suggestions and influence the feature set and interface of the product.

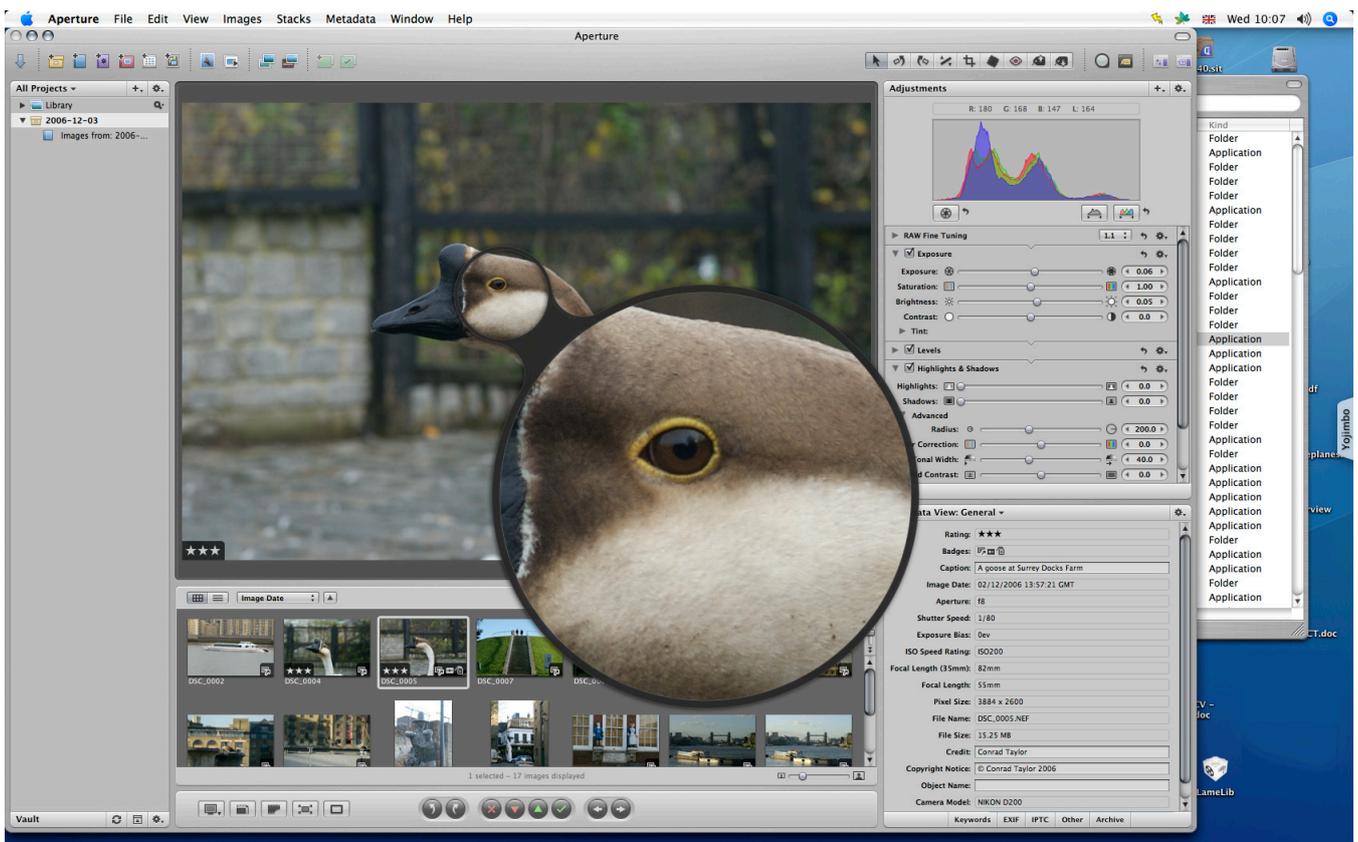


Figure 7. Examining and adjusting a Nikon NEF raw camera image file in Apple Aperture 1.5.

At the bottom is the filmstrip view of the imported images, and the image of the goose is being examined both as a whole and with the Loupe tool for inspecting detail.

In the upper right are the sliders for adjusting how the raw file should be processed, and below that is the zone for adding catalogue metadata.

Aperture and Lightroom are designed to integrate four functions – firstly, they provide useful ‘lightbox’ functionality to help photographers assess the product from a shoot and choose the best shots; secondly, they provide tools for advanced processing of camera raw files; thirdly, they provide image library facilities and means for adding catalogue metadata to assist in image search and retrieval; and fourthly, they provide the means to print image collections as albums for presentation to clients, or convert to a variety of electronic presentation formats such as PDF, Web sites, Flash shows etc.

Raw but standardised?

As we have noted above, one advantage of raw formats could be that as better demosaicing and other image processing algorithms are developed in the future, one might be able to get better results from an archived raw file in ten years’ time that one can now. But photographers may feel nervous about adopting this as an archiving strategy, especially because (a) there are so many proprietary formats (b) most of which are not publicly specified, and may even be encrypted, and (c) camera manufacturers have already shown that they are prepared to turn their backs on digital formats which they invented but in which they have lost interest.

There have been a couple of proposals for an ‘Open RAW’ format, but the standard for raw camera data interchange which would seem to have the best chance of being adopted is the Digital Negative or DGN format put forward by Adobe Systems. This is based on a version of the TIFF standard, and has a published specification.

A small number of digital cameras are now supporting DNG directly, but by releasing the Adobe DNG Converter as a free utility to convert from camera manufacturers' proprietary raw formats to DNG, Adobe has also made it possible to convert, say, from NEF or ORW to DNG, with a reasonable degree of assurance that the converted 'digital negative' will remain accessible for years to come.

One issue which certainly crossed my mind was: if the specifications for the raw files from big companies like Canon and Nikon are proprietary and undisclosed, how can we photographers have confidence that products like Apple's Aperture and Adobe's Lightroom and DNG Converter interpret the contents of those raw files correctly?

The answer, an Adobe representative has assured me, is basically political: it would serve Nikon badly were their products to be seen as posing problems for users of Adobe's popular software. Therefore the camera and big software companies share technical information in an atmosphere of commercial confidence to ensure inter-operability, while still leaving manufacturers of digital cameras the privilege of developing their algorithms in private.

Conrad Taylor

December 2006