

HIGH DYNAMIC RANGE DIGITAL PHOTOGRAPHY

Clipped highlights and shadows are a common problem for digital camera users. Recent research in computer graphics provides a solution, in which a high dynamic range image is constructed from a sequence of different exposures. Guy J Brown FRPS explains




An average exposure of the interior scene (left) shows clipping in both the highlights and shadows. As a result, the window areas are grossly overexposed and there is a loss of shadow detail, as indicated by the histogram (lower left). Above: a high dynamic range image constructed from a sequence of bracketed exposures retains detail throughout the tonal range. The corresponding histogram (right) is not clipped.



At the heart of every digital camera is a sensor that converts incident light into a digital image. A number of factors affect the fidelity with which the original scene is recorded in digital form. Of these, the resolution of the sensor has received the most attention from camera manufacturers, with digital cameras boasting ever-higher megapixel counts.

However, relatively little attention has been paid to the range of values that each pixel may represent. Jpeg images recorded by most digital cameras represent the colour at each pixel with a single byte (8 bits) for each of the red, green and blue colour channels. Although this allows more than 1.6 million different colours to be represented, each colour channel of a pixel can only take



one of 256 possible values. This means that scenes containing very bright and very dark areas cannot be effectively recorded, leading to the familiar problem of clipped highlights and/or shadows. A common example of such a scene is a window-lit interior, as shown on the previous page, below. Note that the corresponding histogram is clipped at both ends; some shadow detail is recorded as pure black, and some highlight as pure white. The situation can be improved to some extent by using the camera's Raw image format, which typically uses 10 or 12 bits per colour channel. However, there are still many circumstances in which the limited dynamic range of a digital camera (typically no more than 8.5 stops) is exceeded by high contrast scenes.

Recent research in computer graphics offers a solution to this problem, in the form of high dynamic range (HDR) imaging. In comparison, the images recorded by conventional digital cameras are regarded as low dynamic range (LDR). A HDR image of the window-lit interior is shown at the top of the previous page; note that detail recorded in the highlight and shadow areas is much improved.

Digital cameras that are able to directly capture the full dynamic range of a high-contrast scene are currently in development, though only a small number of commercial solutions are currently available, and they are either very expensive or tailored to specific applications, such as surveillance video capture.

An alternative to direct HDR capture is to record the scene using a sequence of bracketed LDR exposures, which are then combined using software. The

Above: Construction of a HDR image from five bracketed exposures. Simply scaling the values of the HDR image for output to a screen or printer is not sufficient; the resulting image is too dark and lacks local contrast. Tone mapping gives a much better rendition of the scene.

underlying principle is that each exposure in the sequence will contain some pixels that are properly exposed, and others that are under- or over-exposed. By combining the correctly exposed pixels from each frame and excluding the clipped pixels, a HDR image can be constructed.

If the response of the camera to light were perfectly linear, then each pixel could be divided by its exposure time and averaged (ignoring clipped pixels) to give a HDR image. In practice, however, image sensors do not measure light in a perfectly linear manner, and camera manufacturers apply processing to boost contrast, suppress noise and give a more 'film-like' response. Since the response function of the camera is unknown (and indeed is considered proprietary information by the camera manufacturer) it must be deduced by the HDR software. This may require a separate calibration stage, or an approximate response curve can be used.

HDR software may also compensate for other factors that affect the correspondence between frames of the exposure sequence, such as lens flare, ghosting caused by moving objects and misalignment due to movement of the camera.

Using this approach, HDR image capture therefore involves no more than taking a sequence of bracketed exposures with a conventional digital camera. However, for good results, the following practical considerations should be taken into account: since the technique requires multiple exposures of the same scene, it is most suited to scenes that do not contain significant motion. The best result will be obtained when the camera is supported by a tripod. Hand-holding the

camera is possible – but not recommended – since most HDR software provides the facility to place misaligned images in register. If you have no option but to handhold the camera, use automatic bracketing and continuous drive.

Only vary the exposure time when bracketing. Use manual or aperture priority mode, and keep the ISO constant. White balance should also not be changed during the sequence of exposures. As usual when using a digital camera, optimum quality will be obtained by storing images in Raw, and then converting them to a lossless file format such as Tiff. Also, some HDR software (such as PhotoMatix) is able to construct a HDR image directly from Raw files.

It is advisable to shoot an odd number of bracketed exposures, such that the middle frame has an average exposure for the scene. To meter for the middle exposure, either use the camera's matrix metering, or use an average of spot meter readings taken from the highlights and shadows. Starting from the middle exposure, overexpose and underexpose in steps of two stops, until the camera's histogram display indicates that the shadows and highlights are not clipped. In practice most scenes can be recorded with three or five exposures.

The flow diagram on the previous page illustrates the process of constructing a HDR image from five exposures. In the longest exposure (+4 stops), highlights are severely clipped but the shadows are well recorded. Similarly, the shortest exposure (-4 stops) records full highlight detail, but clips the shadows. By combining the unclipped pixels over the whole sequence of images, the full dynamic range of the scene can be accurately recorded.

Once a HDR image is produced, it must be encoded and stored using an appropriate file format. It should be noted that there are fundamental differences between HDR and LDR image encodings. Each pixel in a LDR image is usually represented by a 24-bit number that describes its position in a colour space, such as Adobe RGB. In contrast, the pixel values in a HDR image are directly related to the radiance in the scene that was photographed.

A number of file formats have been developed that allow HDR images to be encoded in an efficient manner. The most common of these is Radiance RGBE format, which has a .hdr extension. RGBE encoding requires 32 bits per pixel, but allows a far greater dynamic range to be represented than conventional 24-bit RGB encodings.

A far wider range of illumination is captured by a HDR image than can be reproduced by a display device, or in print. Therefore, the problem arises of how to display a HDR image in a manner that is visually acceptable. Simply scaling the contrast range to fit the output device is not sufficient – the resulting image is too dark and bears little resemblance to the original scene. Display of HDR images therefore involves a tone mapping problem: how to preserve detail in the HDR image while reducing its dynamic range to fit the output device. Since the human visual system solves a similar problem, much inspiration for tone mapping algorithms has come from the study of visual perception.

Broadly, tone mapping algorithms can be classified as global or local. Global algorithms compress the dynamic range of a HDR image by operating on each pixel independently, whereas local algorithms adjust the value of each pixel depending on the value of its neighbouring pixels. The latter more closely model the behaviour of the human visual system, and tend to give the more visually pleasing result.

From a photographer's point of view, tone mapping algorithms may be regarded as analogous to methods of contrast control used during wet printing. Global tone mapping algorithms are akin to selecting an appropriate grade of printing paper. Similarly, local tone mapping may be regarded as a means of automatically performing dodging and burning of local areas.

When evaluating HDR software, it is the characteristics of the tone mapping algorithm that will be of most concern to the digital photographer. This point is illustrated in the following two sections, which compare the HDR tools provided by Photoshop CS2 with a specialist application called PhotoMatix. The following descriptions refer to versions of the software running under Mac OS X (versions running under Microsoft Windows are also available).

Support for HDR images was introduced in Photoshop CS2. Select File > Automate > Merge to HDR and choose the images to combine. An option is provided for aligning the images, if required. No separate calibration stage is needed, since Photoshop computes the camera response from the source images.

Controls are then displayed for setting the bit depth and white point of the resulting image. If the 32-bit-per-channel option is chosen, then Photoshop keeps the result as a HDR image, and the selected white point only affects the image preview. A limited number of functions are then available for further manipulation of the image (for example, it can be sharpened but may contain only one layer). It is likely that full support for 32-bit images will be provided in later versions of Photoshop.

If the 8- or 16-bit-per-channel option is selected, then the image values are clipped according to the selected white point, and tone mapping controls are displayed. There are four possible options here: Exposure and gamma – provides controls for adjustment of the overall image brightness and contrast; Highlight compression – compresses the highlight values of the HDR image into a suitable range; Equalise histogram – reduces the dynamic range while attempting to preserve contrast; and Local adaptation – a local tone mapping algorithm which also provides for overall adjustment by a tone curve. Controls are provided for setting the size of the neighbourhood that is taken into account when adjusting the value of each pixel.

The first three of these tone mapping algorithms operate globally, and can be computed rapidly. The equalise histogram algorithm works quite well, but the others rarely give acceptable results. The local adaptation algorithm requires much more computation time, but usually gives the most pleasing image.

In PhotoMatix meanwhile, HDR images are generated by selecting HDR > Generate and choosing the images to combine. Options are provided to use a standard response curve or calculate the camera's response curve from the images (the standard curve usually works well). The images can also be aligned if required. This can either be done automatically, or by a semi-manual method, in which the user specifies control points that mark the same object in different frames.

The HDR image is then computed and displayed with a viewer that allows small sections of the image to be magnified and viewed with an automatic exposure adjustment. A histogram can also be displayed.

The HDR image is then tone mapped by selecting



Result of tone mapping the same HDR image using 'details enhancer' in PhotoMatix (above) and 'local adaption' in Photoshop CS2 (left).

Adobe (Photoshop CS2) www.adobe.com

HDRsoft (PhotoMatix) www.hdrSoft.com

HDR > Tone Mapping. PhotoMatix provides two tone mapping algorithms, a global algorithm called the 'tone compressor' and a local algorithm called the 'details enhancer'. The tone compressor provides controls to determine the overall brightness and contrast of the image, and to set the white and black points. It runs quickly and produces reasonable results. However, most will prefer the look of the details enhancer algorithm. This requires more computation time since it works locally, but often produces beautiful results. The details enhancer also provides more control, allowing the degree of local contrast enhancement to be determined. As before, there are controls for setting the white and black points and overall image brightness, and the colour saturation can also be adjusted.

PhotoMatix also provides other functions to combine shadow and highlight detail from multiple exposures without actually computing a HDR image. The simplest of these is Combine > Average which averages the exposures. Three other 'highlights and shadows' options are provided under the Combine menu, all of which tend to produce more natural results than tone-mapping, and require little or no user intervention. On the downside, they are not as effective at preserving local contrast and may produce rather flat looking images.

PhotoMatix has a comprehensive batch processing facility, that allows the algorithms described above to be applied to a directory containing a large number of images (eg by combining every three images into a HDR image and then tone mapping the result and saving it). The batch processing facility allows HDR images to be computed directly from the Raw files of most camera manufacturers.

The image quality obtained from HDR software is mainly determined by the tone mapping algorithm, and PhotoMatix scores highly over Photoshop CS2 in this regard. The PhotoMatix details enhancer does a better job of rendering shadow detail than the comparable local adaptation algorithm in Photoshop, and produces images with a distinctive ethereal appearance.

PhotoMatix is polished software, backed by a helpful technical support team, and is highly recommended. A further point is that HDRSoft, which markets PhotoMatrix, also markets the tone-mapping part of PhotoMatix in the form of a Photoshop plug-in. This allows its tone-mapping algorithm to be incorporated into a Photoshop workflow.

In the near future, it is likely that HDR imaging will become the norm rather than the exception. Some camera manufacturers (notably Fujifilm) have already introduced sensors with extended dynamic range. However, until HDR cameras become commonplace, the software solutions reviewed here offer an effective means of handling high contrast scenes. In addition, tone-mapped HDR images have an unusual tonality that can be exploited for pictorial effect.

Guy J Brown FRPS
guy@guyjbrown.com

The following text is recommended for those seeking further reading on the subject: High dynamic range imaging: Acquisition, display and image-based lighting by Erik Reinhard, Greg Ward, Sumata Pattanaik and Paul Debevec. Morgan Kaufmann Publishers (ISBN 0-12-585263-0).