

High Dynamic Range Imaging and its Application in Building Research

Axel Jacobs

Abstract

This article describes the theory and application of high dynamic range imaging (HDRI). HDRI is a recent technology allowing for capturing images with a much extended dynamic range whose values represent real-world luminance rather than just arbitrary pixel values. This article lays out the steps necessary to create HDR images, and highlights recent developments in the technology and its applications for building research.

Keywords: high dynamic range; luminance; digital camera; glare

BibTeX:

```
@article{Jacobs07HDR,  
  author = {Jacobs, Axel},  
  title = {{H}igh {D}ynamic {R}ange {I}maging and its  
          {A}pplication in {B}uilding {R}esearch},  
  journal = {{A}dvances in {B}uilding {E}nergy {R}esearch},  
  volume = "1",  
  number = "1",  
  pages = {177--202},  
  year = {2007}  
}
```

INTRODUCTION

Advances in digital camera technology have made traditional wet-film cameras all but obsolete. Consumer digital cameras with resolutions approaching that of film cameras have become affordable and enjoy enormous popularity among hobbyists and professionals alike. This success is mostly due to the instant availability of the images, which also allows for sharing them through the Internet or other digital media. The photographs may be inspected and even modified by image editors before submitting them to online print centres or simply printing them at home. Digital cameras, however, suffer one mayor problem that makes it impossible for them to capture scenes as we see them. While the human visual system can adapt to a dynamic range of up to 10,000:1 for parts of a scene and over 12–14 orders of magnitude in total, cameras have a much lower dynamic range of typically less than 1000:1. As a result, digital images suffer from lack of detail in the shaded regions and overexposed highlights. The pixel values are not correlated to luminance in the real scene, they merely indicate if one object is more or less bright than another.

Recent advances in HDRI have shown how those limitations may be overcome. With HDRI, images of the real world may be accurately captured, stored and displayed. This article reviews some of the technologies used to achieve this goal (Battiato et al, 2003; Reinhard et al, 2006). Apart from having an artistic value, the pixel values of HDR images directly describe the luminance of the objects they depict. With this information and some additional calibration, it is possible to assess the physiological impact of the luminance environment.

CAPTURE

Real scenes encountered in our visual environment usually have large dynamics ranges that may be in excess of 100,000:1 between the brightest and the darkest areas. Digital cameras, however, are only able to capture a rather small part of this. This is due to the properties of the photo sensor within the camera that convert photons to electrons. Due to thermal noise within the active areas, as well as other limitations, a signal will be generated even if no light is incident on the sensor. This dark value limits the lowest irradiance that can be captured by the chip before the usable signal drowns in noise. On the other side of the scale, the sensor can only cope with a certain amount of light before becoming saturated. This not only limits the ability of the one pixel, but the charge may also spill over and effect its neighbours, resulting in the ‘blooming’ of highlights. Figure 1 sketches out the step necessary to convert scene luminance to digital values in a photograph.

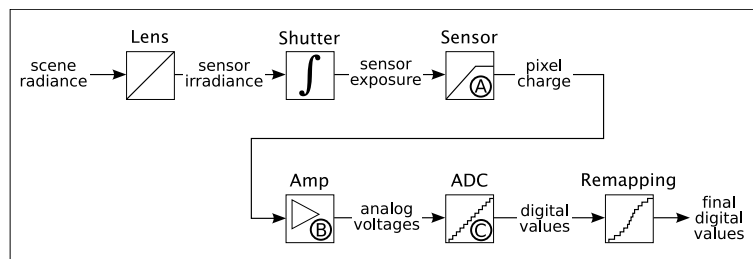


Figure 1: The image acquisition pipeline; *Source:* Adapted from Debevec and Malik, 1997

These two effects restrict the usable dynamic range of the capturing device. As a result, shadows will be entirely black in the image, making it impossible to distinguish any objects in shady areas, while alternatively, highlights above a certain luminance will be overexposed, making them appear just white. An example indoor scene with high contrasts exhibiting those problems is shown in Figure 2.



Figure 2: Igreja Nossa Senhora de Fatima, Lisbon. This is an auto-exposed indoor scene with very high contrasts. Note the lack of detail in the dark areas and the washed out highlights. *Source:* Axel Jacobs

The solution to this problem is to capture the scene with different exposures (Robertson et al, 1999). Many cameras offer an option called auto-exposure bracketing (AEB), typically allowing the capture of five images in intervals of one exposure value (EV). This will create a sequence of two underexposed photographs (−2, −1EV), one autoexposure image (±0EV), and two overexposed ones (+1, +2EV). With a combination of manual exposure correction (MEC) and AEB, this range may be extended to ±4EV, while full manual operating mode is limited

only by the camera's firmware. For example, the fastest shutter speed on the Nikon 995 is 1/2000sec, the slowest one is 8sec. Above this, the camera may be operated in bulb mode, which opens the shutter when the release is pressed until the release is disengaged. Figure 3 is an exposure-bracketed sequence taken in full manual mode.



Figure 3: Sequence of exposure-bracketed image, separated by one f-stop. The shortest shown exposure time is 1/60 sec (top left), the longest 15 sec (bottom right). Aperture fixed at 2.6, ISO 100. Taken with Nikon CoolPix 995. *Source:* Axel Jacobs

A major inconvenience with the built-in AEB feature is often the fact that for each of the typically five exposures, the shutter has to be triggered. For situations like this it is possible to connect the camera to an external control devices, for example the DigiSnap by Harbor (<http://www.harbortronics.com/products2000.asp>) that can be set to take a number of images at once. If more control over the camera is required, or if the AEB feature will not cover the dynamic range of the scene, the only solution for a normal consumer digital camera is to control it through a computer or PDA (personal digital assistant) that may be connected via a USB (universal serial bus) or serial lead, or a wireless interface. A software development kit (SDK) is available, for instance from Canon, allowing developers to design remote control software for most of Canon's cameras. An example of one such software is Automated High Dynamic Range Imaging Acquisition (ADHRIA (http://www2.cs.uh.edu/~somalley/hdri_images.html); O'Malley, 2006). Figure 4 is a schematic showing how the original high scene contrast range may be recomposed from low dynamic range (LDR) photographs into an HDR image.




Figure 4: Inability of a single LDR photograph to capture the huge dynamic range of luminance in our environment. A number of differently exposed LDR images may be combined to reconstruct the whole contrast range in a HDR image. *Source:* Axel Jacobs

Other camera manufacturers might not offer such development kits, or their models might not have remote control enabled. This is true, for instance, for the Nikon CoolPix range of digital consumer cameras. While the professional cameras of the D series can be remotely controlled,

the CoolPix range only offers full remote control over aperture and shutter speeds with older models.

One of the few hand-held HDR cameras is available from TechnoTeam (<http://www.technoteam.de>) and developed in collaboration with Rollei. The camera does not require connection to a computer system when images are taken. Instead, the supplied PC software will combine the bracketed images, applying the factory calibration files along the way. The custom firmware takes one image with an automatic exposure and an additional four with $\pm 1/9^{\text{th}}$ and $\pm 1/3^{\text{rd}}$ of the shutter speed.

Apart from utilizing consumer cameras and for taking exposure-bracketed sequences, camera systems are also available that require control by a computer system. Two such systems are discussed below.

EXPOSURE MERGING

CAMERA RESPONSE

Once a sequence of exposure-bracketed images has been taken, some post-processing is required. The individual images need to be merged, depending on their exposure, into a HDR image (Grossberg and Nayar, 2003a; 2003b; Shaque and Shah, 2004). While it might seem like a straightforward job for a piece of software to add up pixels, it is actually non-trivial. This is because the value of the pixels in a camera's JPEG file is not linearly proportional to the irradiance on the pixel, but rather processed by the camera to give a response curve that is flatter for dark pixels and steeper for more highly exposed pixels. The shape of this curve is generally not available from the camera manufacturer, and may also vary significantly between two cameras of the same make and even model. It therefore needs to be determined by the HDR software.

In a typical imaging system the irradiance E on a pixel is related to the scene luminance L as follows:

$$E = L \frac{\pi d^2}{4 h} \cos^4 \Phi \quad (1)$$

where h is the focal length, d is the diameter of the aperture and Φ is the angle between the optical ray and the optical axis (Mitsunaga and Nayar, 1999). In an ideal system, the intensity recorded by the detector is directly proportional to the exposure time

$$I \approx Et \quad (2)$$

so that

$$I = Lke \quad (3)$$

with $k = \cos^4 \frac{\Phi}{h^2}$ and $e = \pi(\frac{d^2}{4})t$.

e may be described as the exposure that can be altered by varying the aperture or the exposure time. While charge coupled devices (CCDs) are designed to have a linear response over a wide range of exposures, the analog circuitry may introduce non-linearities. Manufacturers also deliberately apply a kind of gamma-mapping to the image to make it more pleasing to the human eye. With M being the final brightness value as stored in the image, knowing the mapping function $g(I)$ allows us to relate pixel value to image brightness.

For this, the inverse mapping function g^{-1} where $I = f(M)$ needs to be determined (Mitsunaga and Nayar, 1999).

Mann and Picard (1995) were the first to point out the problem of response curve calibration and to show a possible solution. They used the function

$$M = \alpha + \beta I^\gamma \quad (4)$$

to model the response curve g . The parameter α describes the dark reading and is estimated by taking an image with the lens covered. The scale factor β is set to an arbitrary value. A regression analysis is applied to determine γ . Although their work is an important milestone on

the quest for radiometric calibration of camera system, Mann and Picard's model does not lead to accurate quantitative results.

Debevec and Malik (1997) propose a method to approximate the camera's response with a logarithmic function

$$\ln f(M_{p;q}) = \ln E_p + \ln T_q \quad (5)$$

with p being the spatial index over the pixels, and q being the index over exposures T . The only restriction is that the exact exposures be known. This model works fairly well for images that are not too noisy.

To overcome the limitations by Debevec and Malik's approach caused by the assumption that the response curve can be described with a logarithmic model, Mitsunaga and Nayar (1999) developed a polynomial approach to the problem. Their approach is also more forgiving in case the image exposure is not known exactly, which is often the case with inexpensive cameras because the repeatability of the aperture setting is limited.

$$I = f(M) = \sum_{n=0}^N c_n M_n \quad (6)$$

The calibration must then determine the minimum order N of the polynomial that is required, as well as the actual coefficients c_n . Response curves for several consumer digital cameras are plotted in Figure 5.




Figure 5: Polynomial response curves for selected digital cameras; *Note:* Only the red channel is shown. The curves are not the result of a thorough calibration, but merely examples from WebHDR. *Source:* Axel Jacobs; data from WebHDR (<http://luminance.londonmet.ac.uk/webhdr>)

The exposure values of the LDR images need to be known for the HDR stitching software to work. They are proportional to the square of the aperture, the equivalent of the ISO rating of the film sensitivity, and inversely proportional to the exposure time

$$EV = \log_2 \left(\frac{f^2}{T} \frac{S}{100} \right) \quad (7)$$

with f being the aperture, T being the shutter speed, and S being the ISO film sensitivity. Fortunately, all this information is stored inside the JPEG files in what is known as the EXIF (exchangeable image file format) header of the image (<http://exif.org>). This EXIF, which is now standard in all digital cameras, stores not only the exposure information, but all settings and parameters of the image and the camera it was taken with. Like with many standards, however, the interpretation of the EXIF standard is often somewhat loose, so the software has to be smart enough to compensate for this. For example, the aperture setting may be stored as *Aperture* or *FNumber*, while *ExposureTime* could also be *ShutterSpeed*, and *ISO* may be *CCDISOSensitivity* with some models.

IMAGE ALIGNMENT

It is evident that the camera's response curve can only be determined if the images are perfectly aligned. For this it is necessary to mount the camera on a tripod and, if possible at all, use a remote trigger. However, because special HDR modes are not implemented in consumer cameras yet, it is usually necessary to manually adjust the settings directly on the camera. This results in camera shake, unless a very heavy-duty tripod is used, which is not practical for capturing scenes other than in a laboratory environment.

While it is possible to align the images, the most common alignment operators are based on edge detection or mean pixel threshold. Those may change considerably between images of the same scene taken with different exposures. To overcome this problem, Ward (2003) proposed the use of an alignment operator based on the median of the pixel values that is fairly robust against changes in exposures. His algorithm generates a median threshold bitmap. An image pyramid is then generated by reducing the size of the image in both directions with every successive step. Each image is aligned with its right-hand neighbour by a maximum shiftwidth of one pixel. The smallest image therefore contributes the most significant bit of the final offset.

This image registration works very reliably for the majority of cases and can successfully align hand-held image sequences. However, the only operation is vertical or horizontal shift. If other operations are necessary such as correction of perspective or rotation, more sophisticated algorithms such as developed by Candocia (2003) or Kim and Pollefeys (2004) need to be applied.

CALIBRATION

RESPONSE CURVE CALIBRATION

Assembling a number of exposure bracketed images into an HDRI involves determining the camera's response function. This automated process is also known as response curve calibration, the result of which is an HDR image. It differs from LDR images in that the pixel information covers a much wider dynamic range. A known response function in combination with the exposure information of the images should in theory also lead to radiometrically correct pixel information. This process is known as radiometric calibration.

PHOTOMETRIC CALIBRATION

The exposure value of the individual LDR image, as determined from the EXIF information of the image, is not always exact for two reasons. First, the aperture is adjusted through a diaphragm inside the lens that has only limited repeatability, introducing an error to the aperture number. Second, the exposure time is not necessarily what is set on the camera. The shutter speed is normally changed in such a way that it is halved with each step as the exposure becomes shorter, however, the value is usually rounded off, for example, 1/4, 1/8, 1/15, 1/30, 1/60. With digital cameras, it is more likely to be 1/4, 1/8, 1/16, 1/32, 1/64 and so on, contrary to what the setting might suggest.

Those errors add up, as a result of which the EV cannot be known very accurately. While this is enough to combine LDR image into usable HDR image, additional photometric calibration must be carried out to obtain per pixel luminance readings of high confidence, as shown in Figure 6.

Calibrated luminance cameras come with a set of calibration files that are the result of laboratory tests of the individual units. For improving the accuracy of HDR photographs, the first and easiest procedure is to compare the HDR luminance with readings taken in the real scene with a spot luminance meter. The ratio of the two may be used as a calibration factor for successive images. Using a Minolta LS-100 luminance meter, the author carried out a rather crude photometric calibration for a Nikon CoolPix 990 and two CoolPix 995 cameras. The resulting calibration factors obtained were 1.11, 1.13 and 1.29 as an average over three readings for a scene under an overcast sky.




Figure 6: The information that can be gained from HDR images and its accuracy depend on the level of calibration; *Source: Axel Jacobs*

Inanici (2006) documents a very thorough evaluation of HDR photography, measuring the relative errors for a large number of coloured and grey targets under different light sources. He used a Nikon CoolPix 5400 with a FC-E9 fisheye lens and derived the following calibration factors: incandescent -0.87 , metal halide -0.86 , overcast -1.9 , and clear sky -2.09 .

Anaokar and Moeck (2005) used achromatic Munsell cards to validate the accuracy of HDR reflectance measurements with various consumer grade digital cameras. Their results demonstrate clearly that cameras with built-in thermal noise suppression performed much better than those without. Another interesting finding is that the luminance measurement of dark surfaces with a reflectance of less than 20% resulted in very large errors. They suggest spectral responsivity calibration to overcome this problem. The authors also looked at relative errors in reflectance measurement for different hues of red, green, blue and purple with low and high chroma. Saturated greens and blues were found to produce the largest errors of up to -80% . The most reliable results could be obtained for warm colours with high Munsell values, while high- to medium-reflectance blues and greens were shown to produce reliable readings down to Munsell value N5. For building interiors, the results should not be too problematic since saturated and dark hues are not frequently found in building materials.

To undertake a more thorough luminance calibration, the vignetting effect has to be taken into account, especially for wide-angle lenses. This was also done by Inanici (2006). His results show that vignetting is virtually zero up to 30° from the optical axis. At the edges of the image, the luminance is reduced by 23% (Goldman and Chen, 2005). See Figure 7 for a plot.

Another systematic error that needs to be corrected is the point spread function (PSF). It is an inherent property of optical systems causing the light passing through the lens to be scattered. This has the effect that the image of an object is somewhat spread out, affecting neighbouring pixels as well. The PSF may be determined computationally by summing up all hot pixels, that is, pixels above a certain threshold in a separate grey scale image (Reinhard et al, 2006). It is then straightforward to remove the PSF from the HDR image by simply subtracting the PSF multiplied by the value of the hot pixel from its neighbours.

SPATIAL CALIBRATION

With a photometric calibration of the camera system, exact photometric information on object luminances may be extracted from the HDR image. If a physiological evaluation of the luminous

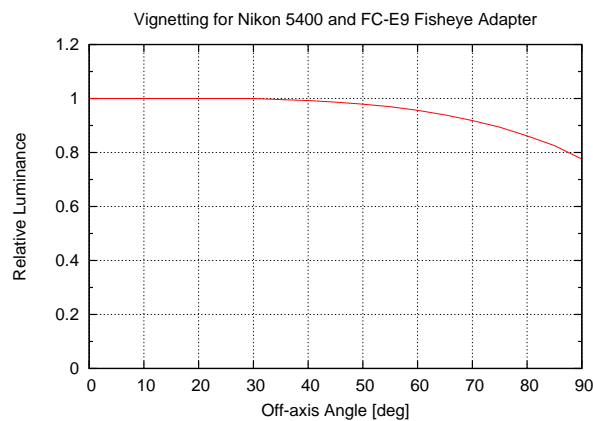


Figure 7: Reduction of image brightness of a fisheye lens image away from the optical axis.
Source: Inanici, 2006

environment is also desirable, for example, glare evaluation or visibility studies, additional spacial or photogrammetrical calibration is required. This process was first described by Wolf et al (1995) and Berutto and Fontoynt (1995). The objective is to determine for each pixel not only the exact object luminance, but also its position with regards to the camera and the solid angle subtended. With this additional information, physiological scene properties such as unified glare rating (UGR) glare indices may be derived. A prototype UGR meter is shown below.

STORAGE

Images from digital cameras are usually stored in 24-bit JPEG format (<http://www.jpeg.org>) with eight bits each for the red, green and blue channels. This results in a resolution of $2^8 = 256$ discrete values for each channel or $2^8 \times 2^8 \times 2^8 = 16$ million values in total. This is also what computer displays can differentiate. Although 16 million colours sounds like a lot, it actually isn't, considering that each channel may only display 256 discrete values. The use of the RAW image format for image capture, which is usually 12 bit or $2^{12} = 4096$ discrete values per channel, may increase the number of discrete values, but to represent the entire dynamic range the human visual system can adapt to, a different file format is needed.

A number of file formats have been developed to overcome the dynamic range limitations of the classical JPEG format as output by digital cameras. They all differ in the dynamic range that they are able to describe, as well as the resulting file size.

RADIANCE RGBE

The most commonly used HDR image format was developed by Greg Ward for the RADIANCE synthetic rendering system (Ward, 1994). Each pixel is represented by a mantissa for each channel and a common exponent. All are stored as 8-bit floating point values, adding up to 32 bits per pixel. The dynamic range of the RGBE format is 76 orders of magnitude (Ward, 1990).

PORTABLE FLOATMAP

The Portable FloatMap HDRI format, which is similar to a floating-point TIFF, assigns 4 bytes to each of the channels. Of those four bytes, one bit is for the sign, one byte for the exponent, and the rest are assigned to the mantissa.

OPENEXR

Developed by Industrial Light & Magic, the OpenEXR file format is now used for all of their film production. It offers a higher dynamic range and colour precision compared to existing 8-bit or 10-bit image formats. The format has support for 16-bit floating-point, 32-bit floating-point,

and 32-bit integer pixels. Multiple lossless image compression algorithms are available, some of which can achieve a 2:1 lossless compression result (<http://www.openexr.com>).

PFS

Although not a file format as such, the PFS format is mentioned here because it introduces an interesting concept for converting HDR file formats that is rather similar to the netpbm tools for LDR images (<http://netpbm.sourceforge.net>). Given a number of N different HDR file formats, a total of N^2 converters would be necessary to allow any format to be converted into any other. With an intermediate format, this number is reduced to $2N$. The PFS stream is designed to be this intermediate format. For this, it needs to be generic enough to hold a variety of data. The PFS format is not meant to be stored on disk or transferred via slow networks (<http://www.mpi-sb.mpg.de/resources/pfstools/>).

Format	Size	%	Quality
RAW image data	4.5 MB	100	Full data
RADIANCE XYZE	1.3 MB	29	Perceptually lossless
8-bit BMP	1.2 MB	27	Loss of HDR
8-bit JPEG, 90% quality	65 kB	1.4	Loss of quality, loss of HDR
JPEG-HDR, 90% quality	70 kB	1.5	Loss of quality, retention of HDR

Table 1: Image size and quality for different encodings; Source: JPEG-HDR White Paper (<http://www.brightsidetech.com>)

JPEG-HDR

It will be apparent that with 32 bits per pixel, an uncompressed HDR image quickly becomes excessively large as image resolution increases. Another problem is that HDR formats are not commonly supported by image editors and viewers. To address this problem, BrightSide developed an HDR extension to the JPEG file format (Ward and Simmons, 2005; Kabaja, 2005).

Table 1 shows the reduction in image size and resulting image quality for a number of different file formats for a given scene.

TONE-MAPPING

Current display technology has a very limited dynamic range, so if an HDR image is to be displayed, the dynamic range of the image will have to be compressed to that of the display device. This process is known as tone-mapping.

According to Reinhard et al (2006), tone-mapping operators can be divided into four categories:

- Global operators – the same (non-linear) curve is applied to all pixels;
- Local operators – the adaptation level is derived for each pixel individually considering the local neighbourhood;
- Frequency domain operators – the dynamic range is reduced based on the spatial frequency of the image region;
- Gradient domain operators – a derivative of the image is modified.

A number of operators exists that try to mimic the human visual system. Human vision is a fairly complex process with a highly non-linear response. Tone-mapping operators require the image to be in real-world units, that is, cd/m^2 . Because the colour-sensitive cones work for high illumination, with the rods being responsible in low-light environments, one of the effects

of human tone-mapping operators is the desaturation of colours for low luminances. Similarly, very bright regions of the image will result in a ceiling luminance thus causing glare (Krawczyk et al, 2005).

The tone-mapping operators briefly introduced below are merely examples chosen because they are implemented in PFSTools (<http://www.mpi-sb.mpg.de/resources/pfstools/>). Many others exist, and the optimum operator for a particular application and output device might require some experimentation from the user (see Reinhard et al (2006) for more exhaustive discussion). Figure 9 is the image of the church again, with different tone-mapping operators applied.

GLOBAL OPERATORS

The simplest possible tone-mapping is linear mapping, but because most display devices exhibit a non-linear response, this will result in very dark images if the HDR image has a wide dynamic range. Gamma-corrected linear mapping attempts to correct this problem. Figure 8 shows an example. Exponential and logarithmic corrections to linear mapping are also possible.

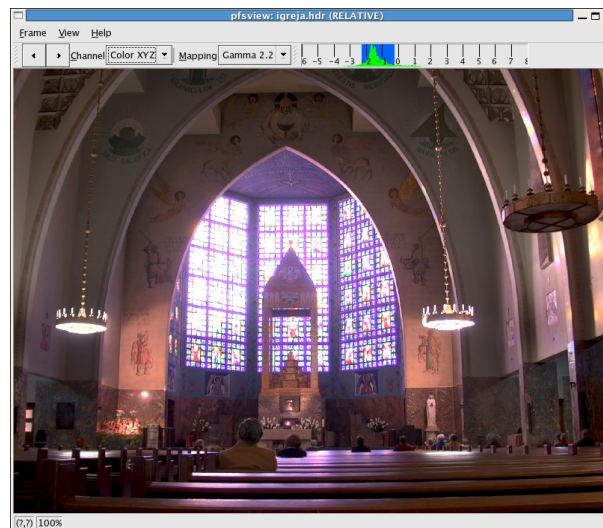


Figure 8: Linearly tonemapped image with gamma correction of the church scene displayed with pfstools; Source: Axel Jacobs

Drago et al (2003a; 2003b) extend the logarithmic response curves in order to handle a wider dynamic range. A logarithmic compression is applied to the image luminance. The base of the logarithm is varied between 2 and 10, based on the brightness of regions within the image. This results in a preservation of contrast in darker regions and a higher compression for bright regions.

Reinhard et al (2002) took inspiration for their tone-mapping operator from techniques known in traditional wet-film photography. To get the visually best print from a negative, it is not enough to just match contrast and brightness. Scene content, image medium, and viewing conditions must often be considered. It was Ansel Adams who developed the so-called zone system in the 1940s. Now, 60 years later, it is still used successfully because it combines quantitative measurements with artistic image content.

LOCAL OPERATORS

The Pattanaik multi-scale observer operator attempts to model all steps within the human visual system currently known well enough to be modelled (Pattanaik et al, 2000). It is the most complete framework to date, although not strictly speaking necessary for only reducing the dynamic range.

In contrast to the Pattanaik model, Ashikhmin's operator only implements those aspects relevant to dynamic range compression. As a result, the Ashikhmin operator is significantly



Figure 9: Results of different tone-mapping operators; *Top left to right*: auto-exposed LDR image as comparison, Ashikhmin, 2002, Drago, 2003a; *bottom left to right*: Durand and Dorsey, 2002; Fattal et al, 2002; Reinhard et al, 2002

faster (Ashikhmin, 2002).

FREQUENCY DOMAIN OPERATORS

It can be argued that an image may be thought of as being composed from a LDR component with a high spatial frequency, and a HDR component for low frequencies. If the two can be separated, only one of the components has to be compressed. This concept has been implemented by Durand and Dorsey (2002). It is based on bilateral filtering, which is an edge-preserving operation that blurs the image while keeping edges intact.

GRADIENT DOMAIN OPERATORS

While high-frequency components within an image cause rapid changes in luminance and high gradients, those are much smoother with low-frequency objects. The gradient field, that is, a representation of the pixel changes, is compressed with operators such as Fattal et al's gradient domain compression (2002). By varying the amount of compression, the amount of detail that is preserved in the displayed image may be adjusted.

DISPLAY

Traditional output devices offer only a limited dynamic range. The maximum contrast of paper prints is fairly low, being only about 100:1. Modern TFT flat panel screens achieve a dynamic range of 300:1. This may be up to 5000:1 for digital video projectors. To make the most of HDR images and display them so neither the dark regions nor the highlights are clipped, new display technology is required. What is claimed to be 'the world's first HDR display' was developed by BrightSide. Their model DR-37P is a 37" high definition HDR display, shown in Figure 10 (Seetzen et al, 2004).

The display achieves a maximum luminance of around 3000cd/m². Combined with the black value of only 0.015cd/m², the overall contrast ratio is as high as 200,000:1. The unit achieves such high luminance and contrast values through the combination of a liquid crystal display (LCD) panel that is not, like conventional flat-panel monitors, uniformly illuminated by cold-cathode fluorescent lamps, but by an array of 1380 light-emitting diodes (LEDs).

The display supports high definition resolution of 1920 × 1080 pixels. Although there seems to be a discrepancy in numbers between the 1380 LEDs and the roughly 2,000,000 pixels of the high definition display. This apparent mismatch is, however, easily explained. While the human visual apparatus can deal with a luminous range of about 1:1000 in any one scene, this range is dramatically reduced for objects that are very close to one another. If there is a fairly dark object close to a bright light source, we will perceive the dark object as black. If the same dark



Figure 10: DR 37P HDR display by BrightSide; Source: BrightSide (<http://www.brightsidetech.com>)

object was put against a background that is less bright, different shades of dark can actually be differentiated.




Figure 11: The output of the BrightSide display is a combination of the LED array (low spatial frequency) and the LCD modulation (high spatial frequency); Source: BrightSide (<http://www.brightsidetech.com>)

It is this effect that the DR 37P exploits. The luminance of the screen is the product of the intensity of the LED and the transmittance of the LCD pixels, as demonstrated in Figure 11. Each are controllable with an 8-bit resolution between 'on' (fully transmitting) and 'off' (not transmitting). In theory, the combined resolution is equivalent to a display with 16 bits per channel or 48 bits in total, however, this is not quite achievable for practical reasons, mostly because the 'off' state of LCDs still allows a not insignificant amount of light to pass through the panel.

APPLICATIONS

LUMINANCE MEASUREMENT

Even without photometric calibration, an HDR image gives a reasonably accurate idea about the luminance distribution in the visual environment and may be used to quickly evaluate scenes such as office environments.

The WebHDR site is very popular with the students on the Masters in Architecture, Energy and Sustainability in the Low Energy Architectural Research Unit (LEARN) (<http://www.learn.londonmet.ac.uk>). It allows the upload of up to nine exposure-bracketed photographs of a scene, which are then processed into a HDR image with Ward's hdrngen (<http://www.anyhere.com>). A screenshot is shown in Figure 12. Tone-mapping and false colours may be applied, and the interactive luminance map can be used to quantitatively inspect the scene. Our students use it a lot, especially for taking photographs inside their scale models. To improve the results, it is also possible to calibrate the camera and reuse those results.

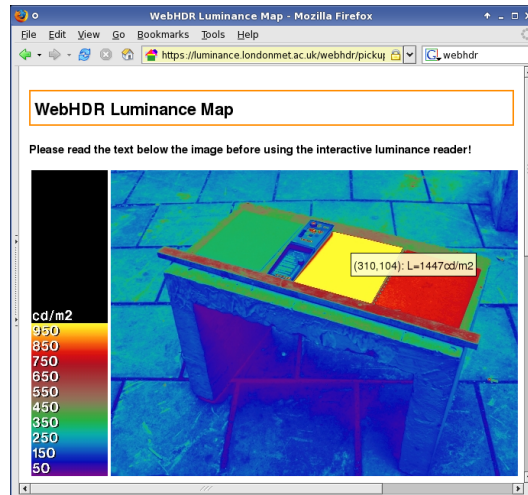


Figure 12: Interactive luminance map in a web browser; *Source:* WebHDR (<http://luminance.londonmet.ac.uk/webhdr>)

If web-based use is not required, there are many other software applications that will create, edit and display HDR images. Some are listed on the WebHDR web site, which may also be a starting point for those wanting to explore HDR photography further.

UGR MEASUREMENT

A small number of HDR-enabled digital capturing systems is already available on the market. Due to the image processing that needs to be applied, and also because of the use of industrial-grade imaging units requiring a frame grabber card, most of these systems are based on portable computer systems, making the set-up and use more difficult than that of an ordinary digital camera.



Figure 13: UGR meter developed by Technical University of Ilmenau; *Source:* TU Ilmenau, Fachbereich Lichttechnik (<http://www.tu-ilmenau.de/site/lichttechnik/index.php?id=843>)

However, because of the computer-based setup, such systems also offer extra flexibility that ordinary cameras cannot deliver. Figure 13, for example, shows a motorized, computer-controlled CCD camera (<http://www.tu-ilmenau.de/site/lichttechnik/index.php?id=843>). Be-

cause both the optical system and the sensor matrix are calibrated, the pixel value may be directly related to object brightness. Additionally, because the spatial characteristics are known from a photogrammetrical calibration, scene properties that are dependent on luminance and object size/location may also be derived. One such application is the measurement of glare, given as UGR value or in any other system (Wolf, 2004). As a reminder, the UGR glare formula is given as

$$UGR = 8 \log \frac{0.25}{L_b} \sum \frac{L^2 \Omega}{p^2} \quad (8)$$

with L_b being the background luminance, Ω being the solid angle of the glare source and p being a position index.

Instead of adding up the glare contribution of every single pixel, the analysis software of the UGR meter combines neighbouring bright pixels into regions, with each region having an average luminance, solid angle, direction and position index. The UGR formula (or any other glare formula that might be applied through the modular system) is then applied to the resulting glare regions.

With photometrically and spatially calibrated HDR images, the exact luminance, direction and solid angle of each pixel is well defined. Only with this knowledge is it possible to derive values such as UGR glare ratings. Other data such as the illuminance are also obtainable. This is implemented in the UGR introduced above for the glare illuminance, that is, the illuminance at the camera for all or individual glare sources (Wolf, 2004). Figure 14 is a screenshot of the UGR meter's control software.




Figure 14: Screenshot of the software of the UGR meter; *Source:* TU Ilmenau, Fachbereich Lichttechnik (<http://www.tu-ilmenau.de/site/lichttechnik/index.php?id=843>)

Moeck and Anaokar (2006) extend this concept further. With templates painted in an image processing program, different parts of an HDR were blacked out. This made it possible to derive the illuminance contribution of different objects to the overall illuminance. It was then possible to look at individual features such as windows, artificial light sources or ceiling.

HDR PANORAMA

Another interesting system is the SpheronCam HDR (<http://www.spheron.com>) that allows for the capture of HDR panoramic (360°) scenes. It offers a dynamic range of 26 f-stops, and can output the panoramas in a number of popular HDR image file formats. The system is depicted in Figure 15. Once set up, the motorized camera can complete an HDR panorama in a few minutes.

No manual stitching or image alignment is required with the Spheron system, all this is done with the software provided. Figure 16 is a sample panorama taken with the SpheroCam HDR.



Figure 15: Panoramic HDR capturing system; Source: Spheron (<http://www.spheron.com>)




Figure 16: HDR panorama of Air Zermatt; Source: Spheron (<http://www.spheron.com>)

For those that don't invest in a SpheroCam HDR, Ward's Photosphere software is capable of assembling HDR panoramas from individual HDR images. It is freely available from Anyhere (<http://www.anyhere.com>), but unfortunately only runs under Mac OS X.

OUTLOOK

SENSOR IMPROVEMENT

HDR is still a fairly specialized field that, although showing great potential, has a long way to go before becoming generally accepted and used by consumers. The main reason for this is that taking HDR images is still relatively complicated and fairly time consuming when compared to 'ordinary' photographs.

In this last section we take a closer look at what happens inside a digital camera and how advances in manufacturing and design of photosensitive chips have the potential to radically change the way digital cameras work, thus enabling them to directly output HDR images without requiring any post-processing.

Part of Figure 1 that shows the photo sensor inside the digital camera is given again in Figure 17. Conversions taking place on the path from irradiance to digital values are marked with A, B and C.

Although it does not qualify as a HDR camera and is only marketed as 'wide dynamic range'




Figure 17: The conversion of light to a digital value in a digital camera; *Source:* Adapted from Debevec and Malik, 1997

Fujifilm's S3 Pro digital camera may be seen as one of the first consumer cameras in this direction. The matrix of CCDs inside the S3 Pro features two types of pixels: S-pixels that do most of the work, and R-pixels that have a much lower sensitivity due to their smaller size. The R-pixels are only used for very bright parts of the image and prevent the highlights from becoming 'blown out'. Fujifilm claim that this approach increases the dynamic range by 400 per cent, as shown in Figure 18.




Figure 18: Extended dynamic range of the Fujifilm SuperCCD sensor; *Source:* FujiFilm (<http://www.fujifilm.com>)

Internally, the S3 Pro utilises an array of CCDs, which is a proven and reliable technology. A major drawback with CCD chips is their relatively slow read-out due to the sequential nature of the process.

Much effort is being undertaken to develop a new breed of imaging sensors based on complementary metal oxide semiconductor (CMOS) technology. The uptake, however, has been relatively slow as a number of problems with this new technology had to be overcome. Products have arrived but are for now restricted to the more expensive range of professional digital SLR cameras. Figures 19 and 20 show the major components in typical CCD and CMOS imaging chips, while Table 2 compares CCD with CMOS in several respects.

CMOS technology differs from CCDs in a number of interesting ways that could potentially lead to a new breed of HDR-enabled cameras (Janesick, 2002; Litwiller, 2005).

With regard to HDRI, the most interesting property of the new breed of CMOS sensors is the fact that they can be read out non-destructively. Combined with the on-chip electronics, this lends itself to the capture of HDR images. Starting with a short exposure time, an automated sequence of images may be taken with increasingly longer exposure times. The on-chip electronics may then combine a HDR image based on the pixel value. Pixels that are too close to the dark value or overexposed would be discarded, while others would contribute towards an exposure-corrected mean value. Due to the integrated on-chip electronics, CMOS sensors are fast enough to even take HDR videos at full frame rates and resolutions (Wandell et al, 1999; Lim and Gamal, 2001).

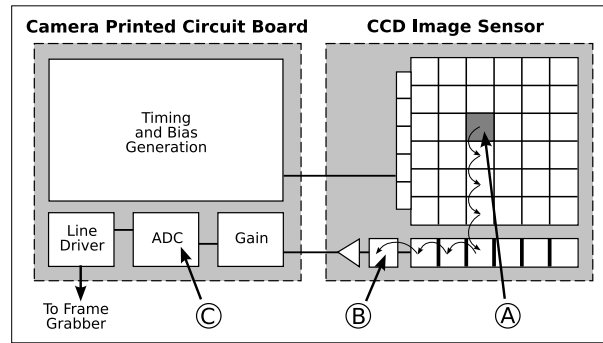


Figure 19: CCD sensor technology. Pixels are read out sequentially; *Source: Litwiller, 2000*




Figure 20: CMOS imaging sensor. Pixel charge is read out individually, resulting in much faster frame rates. *Source: Litwiller, 2001*

FUTURE APPLICATIONS

Sky capture

International meteorological stations monitor weather with automated instrumentation all year round in many locations worldwide. While parameters such as diffuse, global and direct irradiation and illumination are fairly straightforward to determine, sky luminance distributions require relatively costly sky scanners that are made up of a computercontrolled pan-and-tilt head and luminance meter. This setup is limited in its spatial resolution, and a scan typically consists of 145 directions of 11° angles.

HDR images of the sky hemisphere taken with a 180° fisheye lens provide a much finer resolution that is limited only by the number of pixels on the camera sensors. Sophisticated software may then be used to derive parameters such as percentage of cloud cover or correlated colour temperatures. In particular, statistics on cloud cover are of interest to planners and architects interested in optimizing a building's daylight performance. Predominantly sunny skies and predominantly overcast skies required different approaches to get the most from naturally and freely available daylight (Roy et al, 2001; Stumpf et al, 2004).

Glare assessment

While the merits and advantages of one glare rating over another are still being debated amongst lighting experts and physiologists, it is an open secret that daylight glare is virtually impossible to characterize satisfactorily. The daylight glare index (DGI) has been subject of much criticism and is considered by many as practically useless.

The EU-funded project ECCO-Build tried to remedy this situation by deriving a new index, called daylight glare probability (DGP). It is a combination of existing glare indices and an empirical approach based on the vertical illuminance at the eye, the glare source luminance, its solid angle and position index (Wienold and Christofferson, 2005). The measurement of those

Table 2: Comparing CCD and CMOS sensors

CCD		CMOS
Speed	Slow due to sequential read-out	Superior due to on-chip electronics
Windowing	Limited	Possible
Anti-blooming	Susceptible to blooming	Immunity to blooming
Energy consumption	relatively high	low

parameters are based on HDR fisheye photographs. A new software, evalglare, that plugs into the RADIANCE synthetic imaging system was developed to produce DGP values and to visualize the results. Preliminary results show a very high statistical significance and much better repeatability compared to other glare ratings. It is hoped that the DGP after additional validation and with a larger number of , data sets, will become the standard method of assessing glare from windows.

CONCLUSIONS

We have introduced the concept of HDRI. HDR pictures may be produced with any consumer digital camera from a sequence of exposure-bracketed photographs. An HDR image stores values corresponding to the luminance of the real scene with a reasonable accuracy, which may be further improved through additional camera calibration. Although taking HDR images with consumer cameras is time-consuming, the process may be eased with custom HDR systems that are beginning to be available on the market. Applications of HDRI, particularly within the building sector, have been demonstrated.

AUTHOR CONTACT DETAILS

Axel Jacobs, Low Energy Architecture Research Unit, LEARN, London Metropolitan University, United Kingdom;
E-mail: a DOT jacobs AT londonmet DOT ac DOT uk

REFERENCES

- Anaokar S.** and Moeck, M. (2005) 'Validation of high dynamic range imaging to luminance measurement', *Leukos*, vol 2, no 2, pp133–144
- Ashikhmin, M.** (2002) 'A tone mapping algorithm for high contrast images', *Proceedings of 13th Eurographics Workshop on Rendering*, Pisa, Italy
- Battiato, S.**, Castorina, A. and Manucuso, M. (2003) 'High dynamic range imaging for digital still camera: An overview', *Journal of Electronic Imaging*, vol 12, no 3, pp459–469
- Berutto, V.** and Fontoynt, M. (1995) 'Applications of CCD cameras to lighting research: Review and extension to the measurement of glare indices', *Proceedings of 23rd Session of the CIE*, vol 1, New Delhi, pp192–195
- Candocia, F. M.** (2003) 'Simultaneous homographic and comparametric alignment of multiple exposure-adjusted pictures of the same scene', *IEEE Transactions on Image Processing*, vol 12, pp1485–1494
- Debevec, P. E.** and Malik, J. (1997) Recovering high dynamic range radiance maps from photographs', *Proceedings Siggraph*, pp369–378
- Drago, F.**, Myszkowski, K., Annen, T. and Chiba, N. (2003a) 'Adaptive logarithmic mapping for displaying high contrast scenes', *Eurographics*, vol 22, no 3, pp419–426
- Drago, F.**, Martens, W., Myszkowski, K., Chiba, N., Rogowitz, B. and Pappas, T. (2003b) 'Design of a tone mapping operator for high dynamic range images based upon psychophysical evaluation and preference mapping', *Proceedings Human Vision and Electronic Imaging VIII*, HVEI-03, pp321–331

- Durand, F.** and Dorsey, J. (2002) 'Fast bilateral filtering for the display of high-dynamic-range images', *ACM Transactions on Graphics*, Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, San Antonio, TX, pp257–266
- Fattal, R.,** Lischinski, D. and Werman, M. (2002) 'Gradient domain high dynamic range compression', *ACM Siggraph*, Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, San Antonio, TX, pp249–256
- Goldman, D. B.** and Chen, J.-H. (2005) 'Vignette and exposure calibration and compensation', *Proceedings of ICCV '05*, Beijing, China
- Grossberg, M. D.** and Nayar, S. K. (2003a) 'Determining the camera response from images: What is knowable?', *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol 25, no 11, pp1455–1467
- Grossberg, M. D.** and Nayar, S. K. (2003) 'What is the space of camera response functions?', *Proceedings of the 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'03)*, Madison, WI
- Inanici, M. N.** (2006) 'Evaluation of high dynamic range photography as a luminance data acquisition system', *Lighting Research and Technology*, vol 38, no 2, pp123–134
- Janesick, J.** (2002) 'Dueling detectors' Spie's OE magazine, February 2002 Johnson, G. M. and Fairchild, M. D. (2003) 'Rendering HDR images', *IS&T/SID 11th Color Imaging Conference*, Scottsdale
- Kabaja, K.** (2005) 'Storing of high dynamic range images in JPEG/JFIF files', *CESCG*, Central European Seminar on Computer Graphics, Vienna, Austria
- Kim, S. J.** and Pollefeys, M. (2004) 'Radiometric self-alignment of image sequences', *CVPR*, IEEE Proceedings Computer Vision and Pattern Recognition, Washington, DC, pp645–651
- Krawczyk, G.,** Myszkowskiz, K. and Seidel, H.-P. (2005) 'Lightness perception in tone reproduction for high dynamic range images', *Eurographics*, vol 24, no 3, pp635–645
- Lim, S.** and El Gamal, A. (2001) 'Integration of image capture and processing: Beyond single chip digital camera', *Electronic Imaging Conference*, San Jose, CA, pp219–226 Litwiller, D. (2001) 'CCD vs. CMOS: Facts and fiction', *Photonics Spectra*, January
- Litwiller, D.** (2005) 'CMOS vs. CCD: Maturing technologies, maturing markets', *Photonics Spectra*, August
- Mann, S.** and Picard, R. W. (1995) 'On being "undigital" with digital cameras: Extending dynamic range by combining differently exposed pictures', *Proceedings of IS&T 48th Annual Conference*, Society for Imaging Science and Technology Annual Conference, Washington, DC, pp422–428
- Mitsunaga, T.** and Nayar, S. K. (1999) 'Radiometric self calibration', *CVPR*, vol 2, IEEE Proceedings Computer Vision and Pattern Recognition, 1999, Fort Collins, CO, pp374–380
- Moeck, M.** and Anaokar, S. (2006) 'Illuminance analysis from HDR images', *Leukos*, vol 2, no 3, pp211–228
- O'Malley, S. M.** (2006) 'A simple, effective system for automated capture of high dynamic range images', *Proceedings of the Fourth IEEE International Conference on Computer Vision Systems (ICVS 2006)*, New York, p15
- Pattanaik, S. N.,** Tumblin, J., Yee, H. and Greenberg, D. P. (2000) 'Time-dependent visual adaptation for realistic image display', *Proceedings of ACM SIGGRAPH*, vol 34, July, pp47–53
- Reinhard, E.,** Stark, M., Shirley, P. and Ferwerda, J. (2002) 'Photographic tone reproduction for digital images', *ACM Transactions on Graphics*, vol 21, no 3, July, pp267–276
- Reinhard, E.,** Ward, G., Pattanaik, S. and Debevec, P. (2006) *High Dynamic Range Imaging: Acquisition, Display and Imagebased Lighting*, Morgan Kaufmann Series in Computer Graphics, Morgan Kaufmann Publishers, San Francisco, CA
- Robertson, M. A.,** Borman, S. and Stevenson, R. L. (1999) 'Dynamic range improvement through multiple exposures', *IEEE, International Conference on Image Processing*, Kobe, Japan, 23–27 October
- Roy, G.,** Hayman, S. and Julian, W. (2001) 'Sky analysis from CCD images: Cloud cover', *Lighting Research and Technology*, vol 33, no 4, pp211–221
- Seetzen, H.,** Heidrich, W., Stuerzlinger, W., Ward, G., Whitehead, L., Trentacoste, M., Ghosh, A. and Vorozcovs, A. (2004) 'High dynamic range display systems', *Siggraph*, vol 23, no 3, pp760–768
- Shaue, K.** and Shah, M. (2004) 'Estimation of the radiometric response functions of a color camera from differently illuminated images', *International Conference on Image Processing*, Singapore, 24–27 October, pp2339–2968
- Stumpfel, J.,** Jones, A., Wenger, A., Tchou, C., Hawkins, T. and Debevec, P. (2004) 'Direct HDR capture of the sun and sky', *Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa (AFRIGRAPH)*, pp145–149

- Wandell, B.,** Catrysse, P., DiCarlo, J., Yang, D. and Gamal, A. (1999) 'Multiple capture single image architecture with a CMOS sensor', *Proceedings of the International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives*, Society of Multispectral Imaging of Japan, Chiba, Japan, pp11–17
- Ward, G.** [credited as Ward Larson, G.] (1990) 'Radiance file formats', <http://radsite.lbl.gov/radiance/refer/refman.pdf>
- Ward, G.** (1994) 'Real pixels', *Graphics Gems*, vol IV, pp80–83
- Ward, G.** (2003) 'Fast, robust image registration for compositing high dynamic range photographs from handheld exposures', *Journal of Graphics Tools*, vol 8, no 2, pp17–30
- Ward, G.** and Simmons, M. (2005) 'JPEG-HDR: A backwards-compatible, high dynamic range extension to JPEG', *13th Color Imaging Conference*, Boston, MA
- Wienold, J.,** & Christoffersen, J. (2005). 'Towards a new daylight glare rating'. *Lux Europa*, Berlin. September 2005, pp. 157-161
- Wolf, S.** (2004) *Entwicklung und Aufbau eines Leuchtdichte-Analysators zur Messung von Blendungskennzahlen*, Publikationsreihe des Fachgebietes Lichttechnik der TU Ilmenau Nr. 7, Der Andere Verlag, Osnabrueck
- Wolf, S.,** Stefanov, E. and Riemann, M. (1995) 'Image resolved measurement of luminance using a CCD Camera', *Light and Engineering*, vol 3, no 3, pp34–44