



DETERMINING LENS VIGNETTING WITH HDR TECHNIQUES

Axel Jacobs, Mike Wilson

*Low Energy Architecture Research Unit, LEARN
London Metropolitan University, Spring House
40-44 Holloway Road, London, N7 8JL, UK
<a.jacobs@londonmet.ac.uk>*

Abstract

High dynamic range (HDR) photography has seen a tremendous reception by hobbyist and professional photographers alike over the last few years. It has also developed into a cost-effective and flexible tool for the lighting engineer. Pictures taken with almost any digital consumer camera can be exposure-combined into HDR images. With a few simple steps, it is then possible to derive photometrically correct information from those images. This paper will demonstrate how HDR techniques themselves can be used to improve the accuracy to HDR images by determining the vignetting that is exhibited with wide-angle lenses.

Keywords: HDR, photography, camera, lens, vignetting, fisheye, calibration

INTRODUCTION

Traditional wet-film photographic cameras have been made all but obsolete by the advance of digital camera technology. CCD-based imaging devices are now sufficiently small and can be produced at a cost low enough to even implement them as an add-on to appliances that have traditionally nothing or very little to do with photography, such as mobile phones or portable computers and music players.

Although consumer digital cameras suffer from dynamic range limitations more so than film cameras, techniques have been developed allowing to expand their limited dynamic range and to make them usable for fairly accurate luminance measurement. This is of particular interest to lighting engineers as it allows for the assessment of the luminance contribution in the visual environment, and even the evaluation of glare [2].

The algorithms and technologies making all this possible are commonly known as high dynamic range (HDR) imaging or photography. The level of accuracy that is achievable with HDR luminance photographs very much depends on the effort that is put into camera and lens calibration and depends on the intended use. One of effects that might be ignored for normal applications but should be looked into is lens vignetting. This results in the image brightness being lower at the

periphery of the image as compared to the centre. The result is most severe if wide-angle lenses are employed. It is actually possible to use HDR images to accurately measure and quantify lens vignetting. With this knowledge, it is then easily possible to correct the images and eliminate the vignetting.

HIGH DYNAMIC RANGE PHOTOGRAPHY

HDR photography tries to resolve a problem that is exhibited in nearly every digital photograph we take and look at. Between direct sun light and a moonless night, the luminance in our visual environment can have a dynamic range of up to 14 orders of magnitude. While the human visual system can cope with a contrast ratio of up to four orders of magnitude for parts of a scene, the dynamic range of digital cameras is limited to about 100:1. By capturing a scene in multiple images and with different shutter speeds, an exposure-bracketed sequence is created that describes the entire range of luminances occurring in the scene. With the appropriate software which is readily available, those multiple images may be combined into one HDR image. Depending on the intended use of the HDR image, different types of calibration need to be carried out in order to create accurate and reliable results.

Before a set of exposure-bracketed low dynamic range (LDR) images of a scene can be turned into an HDR image, the response function of the system needs to be established. Although the actual CCD chips exhibit a mostly linear response, all digital photographic cameras apply post-processing operations to the raw data. The intended effect is to make the result more pleasing to the human view by having some kind of gamma or logarithmic response. Sadly, the algorithms differ between manufacturers and even models, and even two cameras of the same model might have different response functions.

Different algorithms for approximating the camera's response function have been proposed over the years [3]. The two most widely implemented ones are based on Debevec and Malik's paper [1] and on Robertson [5]. Since the theory of HDR imaging is laid out in other resources, it will not be detailed any further in this paper [4].

An HDR image thus obtained from a set of LDR images contains a much higher range of values than the original LDR images. What's more important, however, is that the information in an HDR is not stored in arbitrary units, but in photometric units, i.e. cd/m^2 instead.

To make this photometric information useful, it is necessary to turn it from relative photometric values into absolute ones by carrying out a luminance calibration. The most straight-forward procedure would be to compare the image luminance to the scene luminance as measured with a spot luminance meter. This will produce a calibration factor which, when applied to the image luminance, will give the correct scene luminance.

This calibration factor, albeit simple and quick to determine, has its limitations when it comes to wide-angle lens systems. All optical lens systems exhibit a property called vignetting which results in the image being darker at the periphery than in the centre. This effect could be ignored for medium-range focal lengths, but becomes very noticeable for wide-angle and ultra-wide angle lenses such as fish-eye lenses.

For medium and long focal length optics, the vignetting can be determined by taking an HDR photograph of an evenly lit wall. The darkening of the pixels with increasing distance from the image centre can then be derived from the image. This

technique can not be used for wide-angle and fish-eye lenses. Such systems would have to be measured in an integrating sphere which is not readily available.

A new approach using a motorised pan-and-tilt head and static light source with very low requirements regarding its luminance uniformity will be presented below.

THE APPROACH

The Camera

Many digital cameras are capable of taking exposure-bracketed images. This is known as auto-exposure bracketing (AEB) and means they can be set to automatically take one auto-exposed frame, as well as an additional one or two images that are over-exposed and under-exposed by up to 1 exposure value (EV). With the typical five-exposure AEB setting, the scene is captured in a sequence ranging from an under-exposure of -2EV to an over-exposure of +2EV.



Figure 1: CoolPix990 with Fish-eye lens

While convenient to use, this range is not enough for capturing scenes with very dark areas, or scenes that have a direct view of light sources or the sky. Those occasions require a wider range of exposures. It is possible to remotely control the camera trigger, as well as some if not all of the camera settings with certain models via USB, serial connection or even a wireless connection. Sadly, this functionality is the exception rather than the norm, and is usually only found in the expensive prosumer or professional range.

The CoolPix 990 as shown in fig. 1 might not be the Nikon's latest model, but it is the last one in the CoolPix range that allows for controlling all exposure options remotely via a USB connection. The camera has a sensor with 3.34 million pixels. The lens is threaded and accepts a number of lens

converters, such as the FC-E8 fish-eye converter. With this lens, the overall focal length of the optical system is 8.0 mm, resulting in a field of view of 183 degrees at the shortest focal length. Since this would have caused large areas of black around a round image, the zoom was set to 19 mm resulting in a full-frame fisheye view, i.e. an image with no black border. The field of view in this setting is 180 degrees across the image diagonal.

It is important to fix the aperture and not change it throughout the entire sequence.

The Light Source

The light source consisted of a cylinder of a 28 cm diameter, 40 cm in length. The inside was painted matte white. The opening was covered with a cardboard mask that has a 15 x 15 cm square cut-out, covered with thin paper. A 16 W compact fluorescent lamps with a 15 cm was mounted inside the cylinder. The paper in the front of the opening created a bright patch sufficiently uniform for the experiments.




Figure 2: False colour image of the light source

The only requirement of the light source and the room in which the procedure is carried out is that the darkest part of the light source is much brighter than the lightest part of the room or laboratory. It is also advantageous if the source has a well-defined cut-off. This can be seen in the false colour representation in fig. 2.

Capture Software

A variety of software for remotely controlling digital cameras in general and CoolPix cameras in particular is freely available on the Internet. None of the ones providing a graphical user interface (GUI) completely satisfied all our requirements, but fortunately a program called *photopc*

(<http://www.math.ualberta.ca/imaging/>) comes in the form of a command-line executable. This is a big advantage since it means the software can be controlled through a script, or, if so required, a custom GUI can be added.

To extend the functionality and improve the ease of use, a graphical interface was written specifically for taking HDR sequences. It was named *photopc control panel (ppccp)* and can be operated in different modes for single shot, sequence, or remote trigger. Images taken may be automatically downloaded. A screenshot of the software is shown in fig. 3.




Figure 3: Screenshot of the PhotoPC Control Panel (*ppccp*)

It is planned to make *ppccp* and its source code available for free download towards the end of 2007.

Camera Mount

In order to determine the vignetting characteristics of the lens, the camera was mounted on a pan-and-tilt head. This device is essentially a motorised table that may be rotated around its horizontal axis, as well as one vertical axis. An image of the head with the camera mounted on it is shown in fig. 4. The device's stepper motors are controlled by computer software. A high gear ratio amplifies the torque and ensures high repeatability.



Figure 4: The pan-and-tilt head

At the begin of the experiment, the pan-and-tilt head was manually moved to the end points. Those are the positions where the light source will appear right in diagonally opposing corners of the image. Once those positions have been set, the controller automatically moved the camera from the start position to the end, stopping at a pre-defined number of step along the way. For this example, 61 steps were set.

When the movement of the camera had stopped, one exposure-bracketed sequence of photographs was taken for each step. Care was taken to ensure that the area of interest, the light source, appears nearly entirely black (under-exposed), as well as completely white (over-exposed) in at least one of the images of each sequence. The shutter speed in each exposure ranged from 1/1000 to 2.0 seconds. Thus, a total of 61 positions times 12 exposures = 732 photographs were taken. The image size was set to 'Full', which on the CoolPix 990 is equivalent to 2048x1536 pixels, and the JPEG compression set to minimal to achieve the best possible quality without going into TIFF mode. This was not possible because the capture time and storage requirements would have been excessively large. With the settings used, a full set of images occupied 750 MB of hard disk space.

Post Processing

As is the case with relatively old digital cameras, the image sensor had developed a number of so-called 'hot pixels' over the years. Those are pixels that give a much higher reading than neighbouring pixels. This is exposure-dependent and increases dramatically with exposures above one second. Unlike its successor, the CoolPix 995, the 990 does not feature automatic noise reduction. Please note that there is a difference between 'hot pixels' as described

above and 'stuck pixels' which always give the full reading, even when no light reaches the sensor. Fortunately, stuck pixels were not a problem with our particular camera.

Hot pixels are determined by taking photographs at different shutter speeds with the lens cap closed. This should result in perfectly black pixels, except for the thermal noise induced by the sensors and the amplifiers.




Figure 5: Effect of exposure time on the number of hot pixels

After careful examination of the influence of shutter speed on the number of noisy pixels (see fig. 5), it was decided to set the threshold to 4%, resulting in 36 hot pixels being removed from the image. This was done by interpolating their value from their neighbours' with a software named *jpegpixi* (<http://www.zero-based.org/software/jpegpixi/>). The diagram in fig. 4 shows how the number of hot pixels relates to exposure time.

The individual image sequences were then combined into HDR images with *hdrgen* (<http://www.anywhere.com>) and *pfshdrcalibrate* (<http://www.mpi-inf.mpg.de/resources/hdr/calibration/pfs.html>). The two different packages were use because they implement different algorithms for response curve calibration. *Pfshdrcalibrate* uses the Robertson [5] approach while *hdrgen* is based on Debevec's paper [1].

Unfortunately, the experiments could not be carried out in a blacked-out lighting laboratory. As a result, some reflections occurred on the glossy paint work in the room. For this reason, the HDR image were thresholded to a value just below the lowest luminance of the light source. The threshold luminance was chosen iteratively so that each light source showed up as a closed polygon

against a black background. Fig. 6 shows all thresholded HDR images superimposed.




Figure 6: Superposition of the individual, thresholded HDR images (inverted).

The pixel intensities, as well as the pixels' x and y positions within the image were extracted for each one of the 61 HDR images. This allowed for the computation of the average brightness of the source and its centre. This data was then normalised so that a pixel that is directly in the corner of the image was given a radius of unit. The pixel in the image centre would have a radius of zero. The values are plotted in fig. 7.

It will be appreciated that the radius can be derived from the pixel's x and y position by mean of the Pythagoras' law, which means that the radius is always positive. The plot in fig 6, however, shows negative values for r if x is less than half of the maximum x position of 2048.




Figure 7: Pixel intensity vs distance from image centre

Please note that no attempt was made to derive an absolute sensor calibration. This would have involved taken sample readings with a luminance meter and would have resulted in a calibration factor to be applied to all pixels. This is a very straight-forward approach, although it should be

carried out separately for different light sources, such as daylight and tungsten.

As can be seen in fig. 7, the two competing programs fared equally well, and there is hardly any difference in the results. The next graph in fig 8 presents a curve-fit of the result obtained with *pfshdrcalibrate*. Different even orders of polynomials were tested.




Figure 8: Curve-fit of the vignetting drop-off (*pfstools*)

- 2nd order: a=-91.7, b=293.5; variance of residuals: $\chi^2=15.5$
- 4th order: c=17.4, d=-105.192, e=294.8; variance of residuals: $\chi^2=14.6$
- 6th order: f=-119.0, g=162.3, h=-148.7, i=296.7; variance of residuals: $\chi^2=12.2$

RESULTS AND DISCUSSION

The fitted curves show a good approximation of the pixel intensities derived from the HDR image. The variance of residuals doesn't decrease much between the different order polynomials, so that the simpler second order function may be used. The effect of vignetting is removed from the image with the following operation:

$$L_{xy} = L_{raw, xy} (1 + 0.31 r^2) \quad (1)$$

with r being the distance between pixel and image centre, normalised to the radius of the circumscribed circle r_{max} .

This may be implemented as a cal file for the Radiance *pcomb* program.

Formula 1 shows that vignetting for this particular set-up (camera, lens, aperture) is 31%. In other words, objects at the periphery of the

HDR image have an intensity (equals scene luminance) of 31% less than the same object in the image centre. If the HDR photograph is applied as a luminance meter for measuring the distribution of luminance in a scene, this must be taken into consideration. A correction of the HDR image, or at least the values and parameters determined from it is an absolute necessity. Since vignetting removal, for instance with the Radiance tool box, can easily be scripted or otherwise automated, it should be easy to accomodate the HDR work flow to automatically or semi-automatically take care of this.

As a side-effect, the two programs *hdrgen* and *pfshdrcalibrate* could be tested against one another and found to both produce repeatable results of the same order. No absolute calibration of lumiance was undertaken since this is not required for vignetting correction.

CONCLUSIONS

It has been shown how HDR photography may be applied to determine the vignetting characteristics of wide-angle and super-wide angle camera lenses. Although it requires taking and post-processing a large number of images, the entire procedure has been automated and can be run with little to no user interference. The inverse of the resulting vignetting function can applied to HDR images thus eliminating the error introduced by lens vignetting.

REFERENCES

- [1] Paul E. Debevec and Jitendra Malik: *Recovering High Dynamic Range Radiance Maps from Photographs*. In SIGGRAPH 97, August 1997
- [2] A. Jacobs: *High dynamic imaging and its application in buildings*, Advances in Building Energy Research, edited by M. Santamouris, James& James, London, 2007
- [3] T. Mitsunaga and S.K. Nayar: *Radiometric Self Calibration*, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Vol.1, pp.374-380, Jun, 1999
- [4] E. Reinhard, G. Ward, S. Pattanaik, P. Debevec: *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, The Morgan Kaufmann Series in Computer Graphics, 2005
- [5] M. Robertson, S. Borman, and R. Stevenson: *Estimation-Theoretic Approach to Dynamic Range Improvement Using Multiple Exposures*, Journal of Electronic Imaging, vol. 12(2), April 2003