

High Dynamic Range Imaging for Computer Generated Imagery applications - a comparison of acquisition techniques

Bachelor Thesis
at the Department
of
Imaging Sciences and Media Technology
University of Applied Sciences Cologne

Pashá Kuliyeu

Matrikelnummer: 11050220

12th February 2009

First Reviewer: Prof. Dr. rer. nat. Manfred Schrey

Second Reviewer: Prof. Dr. Ing. Gregor Fischer

Title: Dynamic Range Imaging for Computer Generated Imagery applications
- a comparison of acquisition techniques

Author: Kuliyeu, Pashá

Reviewers: Prof. Dr. rer. nat. Manfred Schrey / Prof. Dr. Ing. Gregor Fischer

Abstract: In many industrial and advertising applications today three dimensional graphics play a major role in enabling cost effective production workflows. One of the major research areas in the field of computer generated imagery for the last two decades has been photorealistic synthetic lighting of virtual objects. As affordable processing power became available more complex synthetic lighting models could be developed. One of these models derives lighting information of a real-world environment from specially processed photographs. In this thesis, two photographic techniques are presented and compared to the third, commercial one in terms of image quality, acquisition time, image processing time and costs.

Key words: HDRI, CGI, Synthetic lighting, Lightprobe, SpheroCam VR, Fisheye, Mirrored ball

Declaration / Eidesstattliche Erklärung

Hiermit versichere ich, die vorliegende Arbeit selbstständig und unter ausschließlicher Verwendung der angegebenen Literatur und Hilfsmittel erstellt zu haben. Die Arbeit wurde bisher in gleicher oder ähnlicher Form keiner anderen Prüfungsbehörde vorgelegt und auch nicht veröffentlicht.

I hereby confirm that I independently prepared this thesis. No other or auxiliary means than the ones indicated in this work were used for the preparation.
This thesis has neither been presented heretofore in same or similar form to any examination office nor published.

Cologne, 12th February 2009



Unterschrift / Signature

Acknowledgements

I would like to express my gratitude to the following people, whose contributions have made this project possible:

- Dr. Manfred Schrey for supplying me with technical equipment and for his know-how on photographic topics
- The company 809cgi (Johan Visschedijk, Dick Boot, Andreas Rücker) for the continued support, technical advice and equipment
- Dr. Paul Debevec for the invention of the modern HDR production techniques and his valuable advices on this thesis
- Dr. Piotr Świątek for his attendant support and valuable advices on this thesis
- Dipl. Ing. Friedrich Cremer for final printing tasks
- Dipl. Ing. engineer Christian Bloch for his support on the theoretical foundations
- Dave Coffin for his support in the field of HDR imaging and the freeware RAW format converting software
- Geraldine Joffre for her support on the Photomatix® Pro software and her know-how in the field of HDR imaging
- Guido Krebs for providing the Canon® fisheye lens
- Winfried Klarmann for supplying the mirrored ball samples
- My dear friend Thomas Ohle for providing the Canon® zoom lens
- My family, for their continued moral support, without which I would not have been able to complete this project
- Especially my older brother for his practical advice and my fiancée Joanna who helped me to see it through

Introduction.....
Chapter 1 Theoretical foundations	2
1.1 Radiometry	3
1.1.1 Understanding the solid angle	4
1.1.2 Radiant quantities in depth	5
1.2 Photometry	5
1.2.1 Photometric quantities in depth	6
1.3 Color spaces	7
1.4 How the sensor captures the world.....	7
1.4.1 Scene-referred vs. output-referred representation.....	8
1.4.2 Understanding dynamic range.....	9
1.4.3 Exposure values and their relation to the contrast of the scene.....	10
1.4.4 Advantages of working with RAW image data.....	13
1.4.4.1 RAW image data and color spaces	16
1.5 File formats	17
Chapter 2 Mirrored ball acquisition techniques.....	19
2.1 Spherical coordinates and ideal spherical mapping equations.....	20
2.2 Spherical reflections.....	20
2.3 Latitude-Longitude mapping equations.....	21
2.4 Before image capture – things to account for.....	22
2.4.1 Size of the light probe.....	23
2.4.2 Distance between the light probe and the camera.....	25
2.4.3 How many views to capture	26
2.4.4 Focusing on the light probe.....	27
2.4.5 Optical attributes of the light probe	27
2.5 During capture.....	28
2.6 After capture	29
2.6.1 Dealing with lens flares	29
2.7 Determining the reflective attributes of the light probe	30
2.7.1 Purpose of the measurements	30
2.7.2 Basic measurement principle	30
2.7.3 Tested sphere types	32
2.7.4 White balancing.....	35
2.7.5 Accounting for differently sized projections of the paper patch on the balls..	35
2.7.6 Reflectivity measurements and error calculation	35
2.7.7 RAW Conversion.....	36
2.7.8 Preprocessing of acquired images for calculations	36
2.7.9 Calculation results.....	39
2.7.10 Correcting the bracketed light probe image sequence	46
Conclusion.....	46

Chapter 3 Fisheye lens HDRI techniques.....	47
3.1 Ideal fisheye mapping equations.....	48
3.2 Fisheye lens type and practical FOV calculations.....	49
3.3 Image acquisition equipment	51
3.4 Post-Processing Equipment	51
3.5 During Capture – Things to account for.....	52
3.6 Post-processing.....	52
3.7 Filling the Nadir / Zenith point	53
Conclusion.....	55
 Chapter 4 General compositing techniques.....	 56
4.1 Compositing	57
4.2 The multipass rendering workflow	57
4.3 Mathematical background.....	57
4.3.1 The add operator.....	58
4.3.2 The screen operator.....	58
4.3.3 The multiply operator	59
4.3.4 The difference operator	60
4.4 Common rendering passes.....	61
4.4.1 Description of different passes.....	62
4.5 Compositing example.....	64
Conclusion.....	67
 Chapter 5 Comparison of the different HDRI acquisition techniques.....	 69
5.1 The SpheroCam HDR.....	70
5.2 Comparing the techniques	71
5.3 Dealing with light source intensities	76
5.4 A typical HDRI workflow for virtual scene lighting	77
5.5 Final results: Audi R8 CGI renditions.....	78
5.6 Problems of the fisheye technique and a proposal for their solution.....	79
Conclusion and final thoughts	81

Introduction

The computational power of modern hardware made several new synthetic lighting approaches possible, especially in the last decade.

As photorealistic product visualization rapidly becomes a standard in advertising services and industrial applications, service companies are finding themselves in need of standardized workflows due to tight production schedules.

One of the critical areas in computer generated imaging (or images, CGI) is the synthetic lighting and shading of virtual scenes. Often virtual objects have to be combined with photographed scenes and the task of matching virtual and real lighting conditions for proper image integration becomes critical. For years simulated light sources have been used. Major drawbacks of such techniques were elaborate virtual light attribute adjustments and non-physical (i.e. photometric / radiometric) light models.

In 1997-98 Dr. Paul Debevec and his team first presented “a method that [used] measured scene radiance and global illumination in order to add new objects to light-based models with correct lighting. The method [used] a high dynamic range [HDR] image-based model of the scene, rather than synthetic light sources, to illuminate the new objects.” [1] Since then several HDR image acquisition methods for synthetic lighting setup have become popular. In this thesis we discuss two techniques widely used by the CGI industry: mirrored ball and fisheye HDR imaging. We also compare these techniques to a professional solution available on the market in terms of productivity, cost efficiency and image quality.

Each of these techniques can be considered to produce a radiometrically accurate image of the real world. In photography a radiometrically calibrated imaging device is referred to as *light probe*. The light probe images (also called *environment maps* in computer graphics) are used in the process called *Image Based Lighting* (IBL) which derives the lighting information of the real world environment from such an image and uses it as an image-based model to illuminate synthetic objects. At the end of the thesis an IBL-comparison of images acquired with the three different techniques (including the professional solution) is provided.

We hope that this comparative study will help companies in selecting appropriate techniques based on their production demands.

Chapter 1

Theoretical foundations

Before acquiring the light probe image data, we first have to understand how lighting information is stored inside an HDR file and how a real world scene relates to the captured image data. It is also important for computer generated imagery (CGI) as “[...] the measurement of light intensity using precise radiometric (radiance) measurements, or equivalent photometric measurements (luminance), is significant to 3D rendering, particularly with indirect illumination. With physically accurate simulations of light (indirect illumination), radiometric measurements are used to measure the influence of light in the scene, simulating light interaction with surfaces. [...] Rendering is then based on radiometric measurements of light taken within the scene, where light is modeled in accordance to the geometric optics light model. You can see how 3D image generation heavily relies on the physics of light, as well as its measurements.” [2]

The acquisition of high dynamic range images is the process of recording the light intensities available in real world. These intensity values are converted internally by the camera’s sensor to RGB values and are then stored either as raw data, or as pre-processed images on a flash drive. The pre-processed data is considered to have low dynamic range (LDR). These images are directly suitable for printing and displaying purposes on low dynamic range devices, such as monitors or printers. In contrast images stored in a RAW file format are of a higher dynamic range (although they are not “HDR” images in the common sense of this word). They contain the entire information as seen by the sensor during capture and are left untouched by the camera’s internal post-processing routines. As such they can not be directly visualized and have to be processed after the acquisition step.

The fields of science which deal with physics of light and its perception by human vision (and also by digital capturing devices, since these are based on human vision models for image reproduction) are radiometry, photometry and colorimetry.

This chapter provides essential background information on these fields which is relevant to HDR imaging. Specifications of low dynamic range vs. high dynamic range images are given, as well as an overview of the available HDR file formats suitable for light probe images.

1. 1 Radiometry

Radiometry is the scientific discipline that deals with measurements of “optical radiation, which is electromagnetic radiation within the frequency range between 3×10^{11} and 3×10^{16} Hz. This range corresponds to wavelengths between 0.01 and 1000 micrometers (μm), and includes the regions commonly called the ultraviolet, the visible and the infrared.” [3] “The term scene indicates either an artificial or real environment that may become the topic of an image. Such environments contain objects that reflect light.” [4] Hence the notion that everything we see is light – either emitted, reflected, absorbed or scattered by objects – is important for further discussion.

In terms of radiometry light is radiant energy and its unit is joule (J). As light travels through space and different media, different quantities are available to measure several radiometric parameters. Table 1.1 shows the main radiometric quantities and their units.

Quantity	Unit	Definitions
Radiant energy (Q)	J (joule)	Q
Radiant power (Φ)	$\text{J} \cdot \text{s}^{-1} = \text{W}$ (watt)	$\Phi = \frac{dQ}{dt}$
Radiant exitance (M)	$\text{W} \cdot \text{m}^{-2}$	$M = \frac{d\Phi}{dA}$
Radiant intensity (I)	$\text{W} \cdot \text{sr}^{-1}$	$I = \frac{d\Phi}{d\omega}$
Irradiance(E)	$\text{W} \cdot \text{m}^{-2}$	$E = \frac{d\Phi}{dA}$
Radiance (L)	$\text{W} \cdot \text{m}^{-2} \cdot \text{sr}^{-1}$	$L = \frac{d^2 \cdot \Phi}{dA \cdot d\omega \cdot \cos(\theta)} \approx \frac{\Phi}{A \cdot \omega \cdot \cos(\theta)}$

Table 1.1 Radiometric quantities and their units

The radiance approximation holds for small A and ω .

Most radiometric units are derived units. “The base units are a choice of seven well-defined units which by convention are regarded as dimensionally independent: the metre, the kilogram, the second, the ampere, the kelvin, the mole, and the candela. [Base units are defined in *Système International d'Unités* (International System of Units, international abbreviation SI)] Derived units are those formed by combining base units according to the algebraic relations linking the corresponding quantities. The names and symbols of some of the units thus formed can be replaced by special names and symbols which can themselves be used to form expressions and symbols of other derived units.” [5]

1.1.1 Understanding the solid angle

"In radiometry, when considering a point on the surface, many of the measurements are taken with respect to the unit [sphere] surrounding that point. The solid angle is used to refer to some small surface area on the [sphere]." [6]

Plane angles and solid angles are both derived units. When angular data is expressed in radians the following definition found in [7] applies:

"The radian is the plane angle between two radii of a circle that cuts off on the circumference an arc equal in length to the radius."

We can generalize this definition as follows:

$$\theta = \frac{a}{r} \text{ (Eq. 1.1-1)}$$

where θ is the subtended angle in radians, a - the length of the arc, and r is radius. Analogously, the length of the encircled arc is equal to the radius multiplied by the magnitude of the angle in radians.

The solid angle generalizes the plane angle's definition into three-dimensional space and its measure is steradian:

"One steradian (sr) is the solid angle that, having its vertex in the center of a sphere, cuts off an area on the surface of the sphere equal to that of a square with sides of length equal to the radius of the sphere." [7] "The solid angle is thus [the] ratio of the spherical area to the square of the radius." [3] The equation is given by

$$\Omega = \frac{S}{r^2} \text{ (Eq. 1.1-2)}$$

where Ω is the solid angle in steradians, S is the surface area of a part of a sphere and r is the radius of the sphere.

Figure 1.1 shows both plane and solid angles. Please note that the solid angle can only be illustrated schematically, because "the solid angle is the quantitative aspect of the conical slice of space, that has the center of the sphere as its peak, the area on the surface of the sphere as one of its spherical cross sections, and extends to infinity." [8] Hence, it can not be visualized directly.




Figure 1.1 Radian and steradian units in comparison

For the sake of clarity only a hemispheric view is given for steradian unit depiction.

If we analyze the dimensionality for both angle types we find that:

$$[\theta] = \frac{\text{m}}{\text{m}} = 1 \quad \text{and} \quad [\Omega] = \frac{\text{m}^2}{\text{m}^2} = 1 \quad (\text{Eq. 1.1-3})$$

Thus plane angles and solid angles are dimensionless quantities. In practice the unit symbol “sr” is used to distinguish between plane and solid angles.

1.1.2 Radiant quantities in depth

As mentioned before light is radiant energy measured in joules.

Radiant power (also called radiant flux) Φ is the derivative of energy with respect to time. The power may be the total emitted from a light source, or the total arriving at a particular surface, and its unit is joules per second, or watt.

Radiant power can further be measured per unit area, which is **irradiance** E (also called radiant flux density). “Irradiance is power per unit area incident from all directions in a hemisphere onto a surface that coincides with the base of that hemisphere.”[3] While irradiance is a measurement of the amount of light arriving at a point from all directions, **radiance** L quantifies the amount of light arriving at a point from a specific direction, per unit area.[6] A similar quantity to irradiance is **radiant exitance** M , “[...] which is power per unit area leaving a surface into a hemisphere [into all directions] whose base is that surface.”[3]

Radiant intensity I describes light emitted into a specific direction. It is the measurement of energy per unit of time per unit solid angle. The integral of radiant intensity over solid angle is power.

1.2 Photometry

Rather than dealing with the absolute quantity of light (that is the amount of light available in the scene), the field of photometry deals with brightness levels as perceived by the human vision. Photometric quantities are in essence radiometric quantities “[...] weighted for the spectral response of the human eye [...]”[3]. As such the quantity symbols stay the same except the subscript “v” added to them.

Quantity	Unit
Luminous energy (Q_V)	lm s
Luminous power (Φ_V)	lm (lumen)
Luminous exitance (M_V)	lm · m ⁻² = lux(lx)
Luminous intensity (I_V)	lm · sr ⁻¹ = cd(candela)
Illuminance (E_V)	lm · m ⁻² = lux(lx)
Luminance (L_V)	cd · m ⁻² = nit

Table 1.2 Photometric quantities and their units

Radiometric units are converted to photometric units by weighting them with the **luminous efficiency function**, $V(\lambda)$. Human vision normally only responds to wavelengths of “[...] approximately 380 to 830 nanometers (nm). Within this range, the human eye is not equally sensitive to all wavelengths. In addition, there are differences in sensitivity to the spectral composition of light among individuals. However, this range of sensitivity is small enough that the spectral sensitivity of any human observer with normal vision may be approximated with a single curve. Such a curve is standardized by the Commission Internationale de l’Eclairage (CIE) and is known as the $V(\lambda)$ curve (pronounced vee-lambda) [...]” [9]




Figure 1.2 The $V(\lambda)$ luminous efficiency function

Human vision reaches its maximum sensitivity at about 555 nm.

1.2.1 Photometric quantities in depth

Below several important photometric quantities are discussed.

Luminous intensity is an important photometric unit. Its unit - the candela, is one of the seven base units of the SI system, and as it is defined as the luminous intensity, in a given

direction, of a source that emits monochromatic radiation of frequency 540×10^{12} hertz and that has a radiant intensity in that direction of 1/683 watt per steradian.

The ***lumen*** is an SI derived unit for luminous flux, it is derived from the candela and is the luminous flux emitted into unit solid angle by an isotropic point source having a luminous intensity of 1 candela.

“***Illuminance*** is another SI derived quantity which denotes luminous flux density. It has a special name, lux, and is lumens per square metre, or lm/m^2 . [...] Most light meters measure this quantity, as it is of great importance in illuminating engineering. The IESNA Lighting Handbook has some sixteen pages of recommended illuminances for various activities and locales, ranging from morgues to museums. Typical values range from 100 000 lx for direct sunlight to 20-50 lx for hospital corridors at night.”[3]

“***Luminance*** should probably be included on the official list of derived SI quantities, but is not. It is analogous to radiance, differentiating the lumen with respect to both area and direction. It also has a special name, nit, and is cd/m^2 [...] It is most often used to characterize the ‘brightness’ of flat emitting or reflecting surfaces. A typical use would be the luminance of [a] laptop computer screen. They have between 100 and 250 nits, and the sunlight readable ones have more than 1000 nits. Typical CRT monitors have between 50 and 125 nits.”[3]

1.3 Color spaces

The CIE XYZ color space was derived from color perception experiments and encompasses the whole range of color values humans normally perceive. This color space is used as a reference for device dependent spaces which are related to it by a set of ***color primaries***. For any given system (human vision, LCD displays, printers) color primaries define a certain maximum amount of colors the system is limited to - the ***color gamut***. Color primaries are device-dependent: while display devices use additive color mixing (mainly red, green and blue are used as primaries), printing is based upon the subtractive method (typically based on cyan, magenta, yellow and black). While mixing color primaries in the additive mode with their maximum intensity yields white, mixing primaries in the subtractive mode results in black.

The overall relationship between perceived and recorded color information is described in [11]: “Humans can see about 10 to 13 million colors. Digital imaging with 24-bit color defines approximately 17 million colors.”

1.4 How the sensor captures the world

In the following discussion we assume that a digital capture device (camera) with the sensor as photographic medium is used. We also assume HDR assembly from digital images, although analog images are suitable as well.

1.4.1 Scene-referred vs. output-referred representation

When images are recorded by a camera they are generally altered by the camera's internal routines to meet the requirements of devices they are to be displayed on, unless a RAW format is used for capture.

Thus these images are identified as “*output-referred*”. We can distinguish between tonal adjustments and additional filtering adjustments (like camera sharpening and noise reduction) performed by the camera's internal processing. While filtering is hard to identify as a particular transformation (since this kind of information is mostly proprietary to camera manufacturers), tonal transformations are already provided by color space transformations, which are applied to raw sensor image data when it is saved to an output-referred format. As stated in [12] “Most image encodings fall into a class called output-referred standards, meaning that they employ a color space corresponding to a particular output device rather than to the original scene they are meant to represent. The advantage of such a standard is that it does not require any manipulation prior to display on a targeted device, and it does not “waste” resources on colors that are out of this device gamut. Conversely, the disadvantage of such a standard is that it can not represent colors that may be displayable on other output devices or that may be useful in image processing operations along the way.” Thus color values recorded by the sensor prior to the transformation may get clipped through the color transformation, which may result in a relatively limited gamut. Additional color values may be useful in HDR processing such as accurate white balancing or better color sampling during IBL. Furthermore as image data acquired by the sensor gets compressed for display purposes (raw data can not be displayed directly) the compression can result in additional quality loss, as camera manufacturers most frequently use JPEG compression. Mainly its *lossy* variant is used, which trades off image quality for smaller image files. Particularly fine details can get lost, since the quantization algorithm used by the JPEG compression reduces high frequency content first.

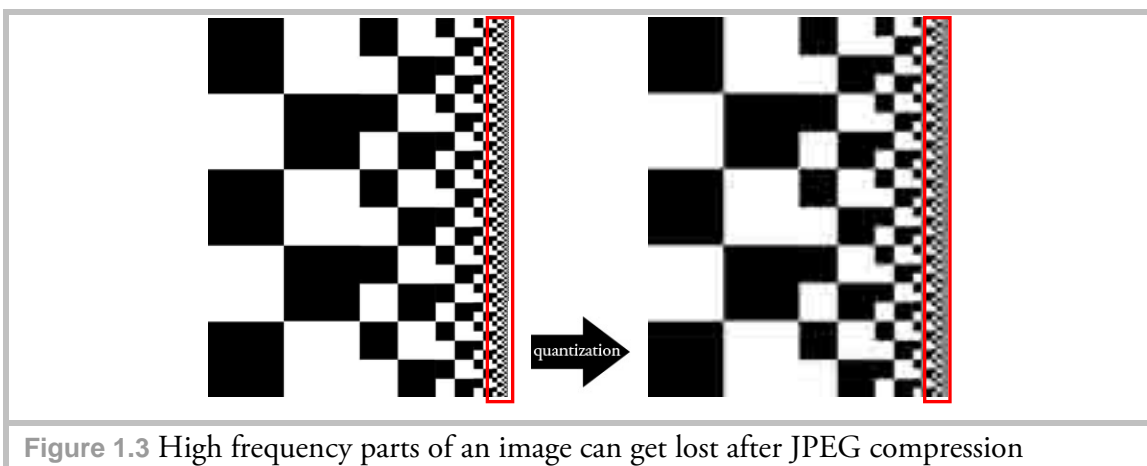


Figure 1.3 High frequency parts of an image can get lost after JPEG compression

In contrast to the output-referred imagery the *scene-referred* image acquisition“ [...] follows a different philosophy, which is to represent the original captured scene values as closely as possible.“ [12] As mentioned before the world is perceived by the human vision as varying light levels, hence a scene-referred approach records the lighting information of the scene. We can consider every digital camera to be a radiometric measurement device

as voltage levels captured by the sensor correlate to radiance levels in the scene. Hence it is of great importance that this information be passed unaltered to the HDR assembly process. As RAW images retain the linearity of the sensor response it can be partly considered a scene-referred standard, although the DR captured is very small compared to that of the scene.

In summary the relevant characteristics of image acquisition are as follows:

	Scene-referred image acquisition	Output-referred image acquisition
Purpose	<ul style="list-style-type: none"> ▪ Main goal: Capture light information of a real world scene ▪ HDR imaging domain 	<ul style="list-style-type: none"> ▪ Main goal: capture images for direct display and photographic processing ▪ LDR imaging domain
Image bit depth	32 bits per channel	8 bits / 16 bits per channel
Data type	Floating point	Integer
Container formats to store data	HDR (discussed later) OpenEXR (discussed later) RAW (can be considered partly an HDR format)	JPEG TIFF and any other format that can hold compressed / uncompressed integer data
Application areas	Radiometric measurements, image analysis, photography, CGI applications	Photography

1.4.2 Understanding dynamic range

In photographic terms (as there are other applications too) the *dynamic range* (DR) describes the ratio between the maximum and minimum light intensities (luminance) in a scene. DR is directly related to the term “*contrast ratio*” which describes more generally the ratio between any sort of darkest and brightest values.

As mentioned at the beginning of this chapter object surfaces react to incident light by either scattering, transmitting, absorbing or reflecting it. Because of this behavior no pure black or white levels (as far as radiometric / photometric measurements are concerned) are ever present in a scene. Therefore to properly determine the DR of a scene two parameters are required: the overall range of brightness and the smallest step [13]. The smallest step is the lowest light intensity found in a scene which is still distinguishable from pure black value. For instance, if the maximum light intensity of a scene is 100 cd/m² and the darkest shadow is 1 cd/m² luminous, then the DR is 100:1.

DR is a device / medium dependent ratio. As light information gets translated from the scene through the capturing device (camera) to a reproduction medium (display, paper) the DR continuously decreases. The field of imaging concerned with translation of DR values in such a way that best parts of the original DR remain preserved is called *tone mapping*. Figure 1.4 shows an example of DR reduction.

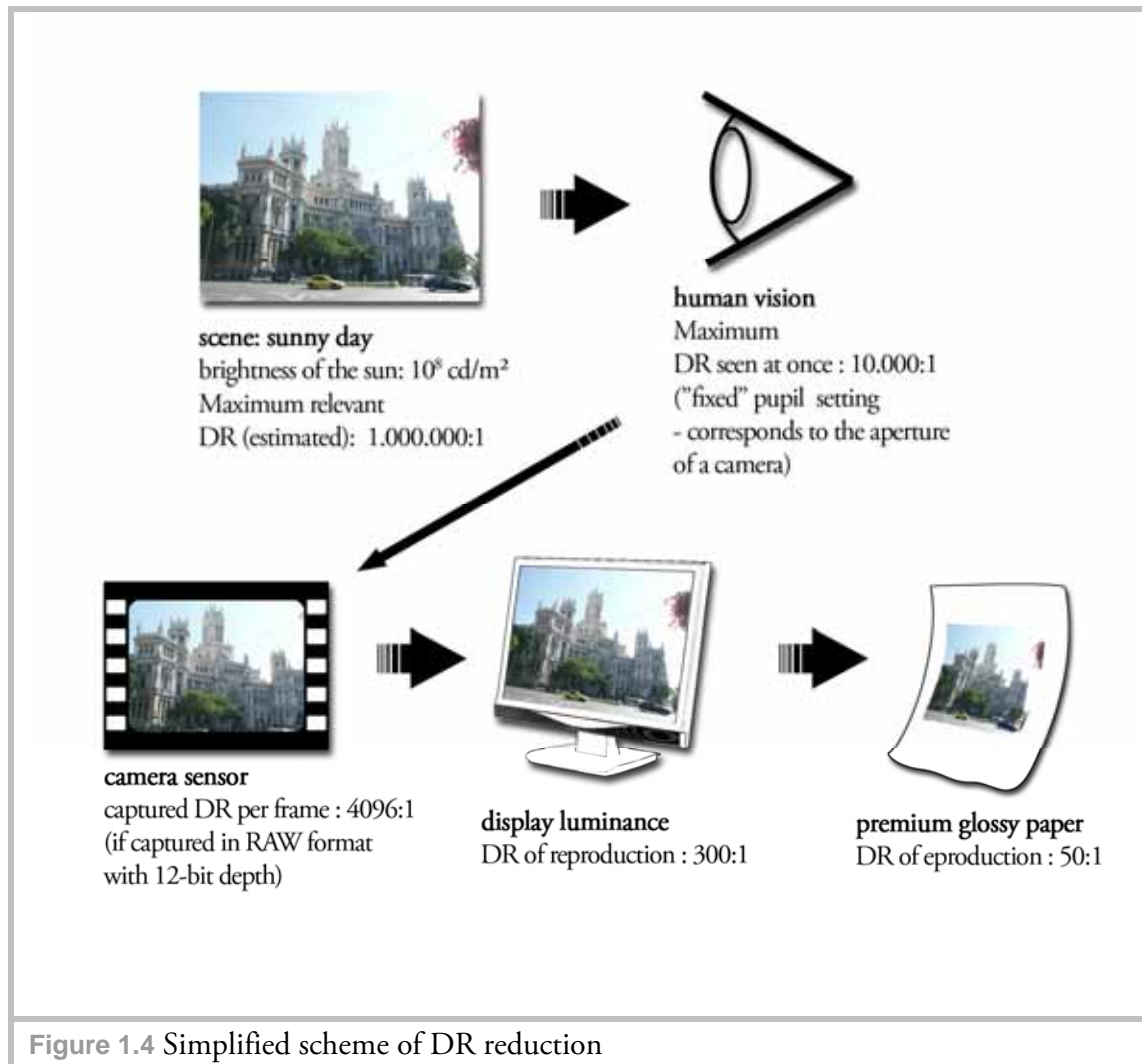


Figure 1.4 Simplified scheme of DR reduction

1.4.3 Exposure values and their relation to the contrast of the scene

The term *exposure* describes the amount of light which hits the sensor and is defined by the following factors: aperture (defined by the "f-stop" value in photographic jargon), exposure time and sensor sensitivity. The same exposure can be achieved by many different combinations of these factors. The generic exposure equation based on the averaged luminance of the scene is given by:

$$\frac{t}{N^2} = \frac{K}{L_s S} \quad (\text{Eq. 1.4-1})$$

where:

t	exposure time [seconds]
N	relative aperture ["f-stop" number]
L_s	metered average scene luminance [cd/m^2 , or some other appropriate unit]
S	sensitivity (speed) of the film or digital images [ISO speed number]
K	an arbitrary constant, which accommodates the units involved and reflects a particular manufacturer's concept of "proper exposure".
	K is often spoken of as the "reflected light exposure meter calibration constant"
	[14]

In photometric terms, exposure is the total amount of illuminance per unit area of the sensor per unit of time and its measure is lux per second.

Exposure value (EV) is a logarithmic scale parameter which can be derived by applying a base-2 logarithm operator to both sides of the exposure equation. When applied to the left-hand side, EV specifies combinations of exposure time and relative aperture which yield the same exposure (rounded values are used). An example is given in Figure 1.5.

$$EV = \log_2 \frac{N^2}{t} = -\log_2 \frac{t}{N^2} \quad (\text{Eq. 1.4-2})$$

e.g.: $-\log_2 \frac{1/500\text{s}}{(2.8)^2} \approx -\log_2 \frac{1/250\text{s}}{(4)^2} \approx 12 \text{ EV}$

relative f-stops	2.8	4	5.6	8	11	... =12 EV
seconds	$\frac{1}{500}$	$\frac{1}{250}$	$\frac{1}{125}$	$\frac{1}{60}$	$\frac{1}{30}$	


Figure 1.5 Exposure calculation example

Different aperture and exposure time combinations lead to the same EV.

The above definition of EV is an absolute definition based on the standardized EV scale defined by the ISO (International Organization for Standardization). Here the base EV (EV 0) corresponds the exposure time of 1 second at a relative aperture of 1.0 for a film (or sensor) sensitivity of ISO 100.

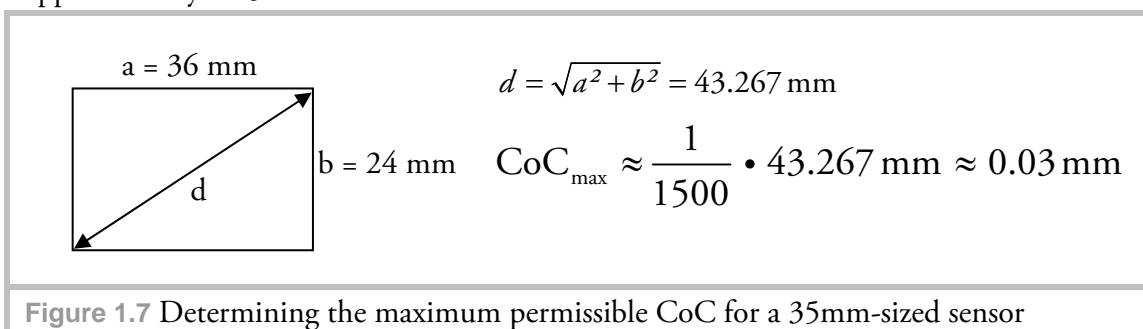
In practice EVs define a relative unit of the photographic exposure scale. Photographers often refer to EVs as "f-stops" which is not an accurate description since both exposure time and aperture define EVs. In this respect each increment of 1 EV corresponds to half as much light reaching the sensor, either by altering the exposure time or the aperture setting.

To determine the amount of EVs needed to cover the DR of a scene spot metering is used. For that either the exposure time or the aperture setting for the darkest and brightest regions of the scene is measured. DR of the scene in terms of required EVs can then be calculated (an example is shown in Figure 1.6).



In photography both modes “*aperture priority*” (AP) and “*shutter priority*” (SP) (i.e. exposure time priority) are available. While AP varies the amount of light by altering exposure time at a fixed aperture, SP does the opposite. If a variable aperture is used the aperture size “[...] not only determines the amount of light that reaches the sensor but also changes the amount of acceptable focus of the image. The *depth of field* (DOF) is the range of acceptable focus from close to the camera to the most distant within the scene [...]”[15] Distance corresponding to the DOF on sensor side of the lens is called *depth of focus*.

“Each point of light in the image focuses a cone of light, with the perfect focus making a circle of light on the sensor surface.”[15] If we consider an infinitesimal point of light being projected on the sensor through the lens, rays shot from this point toward the sensor are never focused on a perfect point as this would demand a zero sized projection. Thus all projections on the sensor are defocused to a certain extent and are therefore called “*circles of confusion*” (CoC). An acceptable size of a CoC which is still perceived as being in focus generally depends on the resolution of human vision and the viewing distance at which the acquired image is observed, and is known as the *maximum permissible CoC* (CoC_{max}). A widely used CoC_{max} is approximated to be 1/1500 of the sensor’s diagonal, which corresponds to resolving 5 lines per millimeter on a print of 30 cm diagonal. If we assume our sensor to have the dimensions of a 35mm film (several high-end digital SLR cameras are equipped with such sensors) this leads to a CoC_{max} of approximately 0.03 mm:



Each CoC whose diameter is smaller (or equals) the CoC_{max} will appear to be in focus in the resulting image. The next figure illustrates the above discussion (optical paths have been simplified for demonstration purposes).

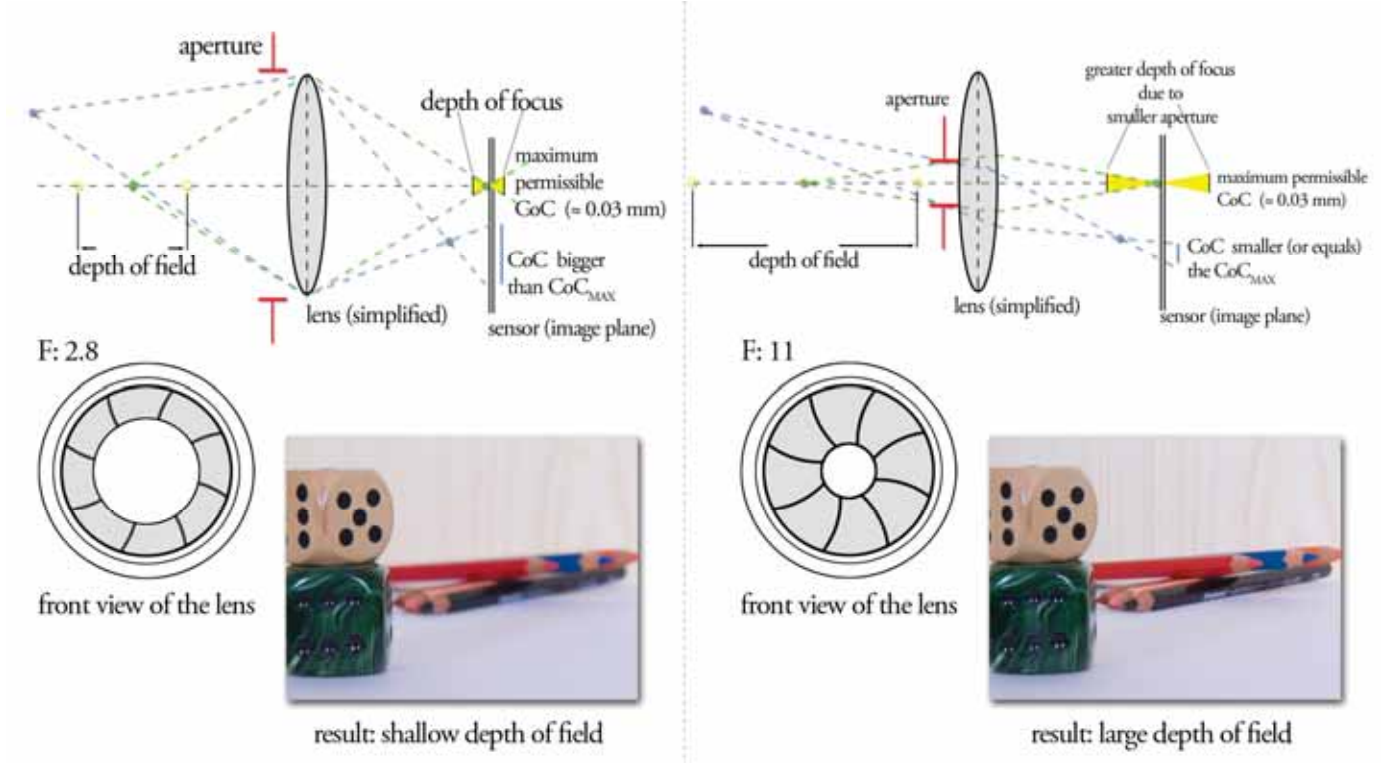


Figure 1.8 In the left image the green point is completely in focus, while the blue point produces a defocused spot on the sensor. In the right image a smaller aperture is used so both green and blue points are in focus. Resulting image on the right was exposure-time compensated for the smaller aperture size to retain the same exposure for demonstration.

Hence if the HDR image were to be assembled from an aperture bracketed image sequence the DOF would vary from one frame to another, which is undesirable. Therefore varying the exposure time at a fixed aperture setting in the SP-mode is the proper way to acquire bracketed images for HDR assembly.

The conversion between EVs and the DR of a scene is given by:

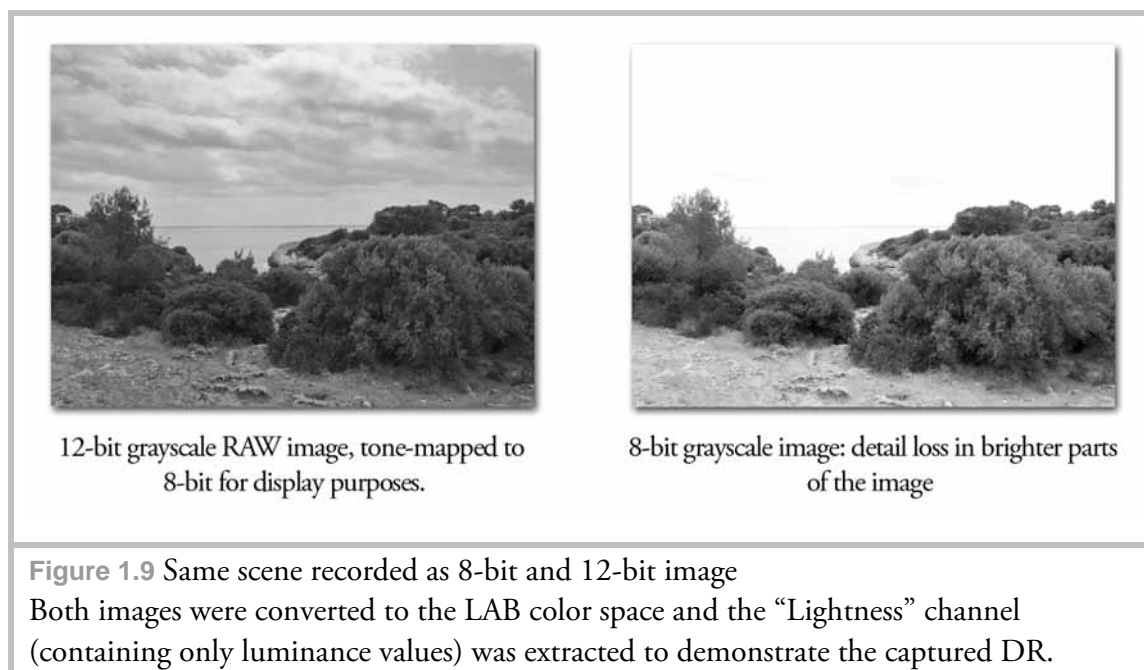
$$2^{\text{EV}} = \text{DR} \text{ (Eq. 1.4-3) [16]}$$

For instance a scene whose DR is 1.000.000:1 would require $\log_2(10^6) \approx 20$ EVs to cover the entire DR.

1.4.4 Advantages of working with RAW image data

HDR images can be assembled either from pre-processed compressed image files (i.e. JPEGs or TIFFs) or from RAW image data which has to be manually converted for

HDR assembly. Both these ways of storing lighting information can be considered to produce LDR images even though RAW data with its 12-bit depth per channel allows for more fine-grained capture of light levels. This provides an advantage in terms of the number of exposures that have to be taken to cover the DR of a scene: as each image contains $2^{12}=4096$ levels of gray, fewer images are necessary.



As described in section 1.4.1 images recorded in a compressed format can be considered output-referred as they are preprocessed for display purposes. As non-linear images do not provide information on how pixel values relate to the luminance of the scene, the camera response curve has to be generated so that the non-linear color transformation can be reversed. This curve is responsible only for in-camera processing steps and should not be confused with the sensor curve, which is generally linear over its entire dynamic range. Calibration techniques and additional information can be found in [17]. HDRShop, a software package developed by Dr. Paul Debevec (<http://www.hdrshop.com/>), has an integrated option to calibrate the response curve from a reference bracketed image sequence.

As images are processed after capture by the camera they eventually get transformed into one of the standard color spaces. Hence an assumption can be made that the inverse transformation will linearize the image and in fact some software packages provide such an option. Nevertheless in general this approach is not recommended, as camera manufacturers apply additional color adjustments (such as contrast adjustments) to achieve a more pleasing result and consequently alter the non-linear color transformation [18].

On the other hand RAW image data is the closest representation of what the sensor receives during image acquisition. At the heart of almost all digital photography today is the sensor of type known as ‘mosaic sensor’ or ‘color filter array’ (CFA) [19] “Color filter array cameras use a two-dimensional area array to collect the photons that are recorded in the image. The array is made up of rows and columns of photosensitive detectors - typically using either CCD (charge-coupled device) or CMOS (complementary metal

oxide semiconductor) technology - to form the image. In a typical setup, each element of the array contributes one pixel to the final image [...]”[19] Such a sensor generates electrical charges proportional to the luminance of a scene. Thus the image produced contains only grayscale intensity values. For color acquisition a color filter array is added. “Each element in the array is covered by a color filter, so that each element captures only red, green, or blue light.”[19] Since each element captures only one of the primary colors, portions of an image formed by groups of elements of the same color can be represented as grayscale intensities of that particular color.

After recording, RAW data has to be converted for further presentation or editing. As the layout of array elements and array’s architecture itself vary among different manufacturers, *decoding* and *demosaicing* information differs and is stored by manufacturers in their proprietary RAW file formats as metadata. Decoding phase “[...] tells raw converters which color each pixel [array element] represents.”[20] Demosaicing is the interpolation process of “[...] the ‘missing’ color information for each pixel from its neighbors.”[20] Bayer pattern color filter arrays are widely used in consumer electronics. At Siggraph 2005 an explicit overview of the benefits of direct RAW data processing was presented [21]. For the purpose of the review the authors used unprocessed data from several digital cameras (including Canon® 1D Mark II, the predecessor of Canon® 1Ds Mark II used for tests in this thesis). As mentioned before the response of sensors in modern digital cameras can be assumed to be linear relative to the input luminance of the scene. The authors photographed an HDR scene comprising 36 exposures (at 1/3 EV step size). For four levels of luminance the recorded raw pixel values (after decoding and demosaicing) were tested for linear relation to exposure times at which the individual images were taken. The test showed linear consistency of the sensor data across the whole recordable dynamic range. As light intensities in the scene are linearly proportional to the recorded voltage levels the image acquisition can thus be considered radiometrically correct.

Figure 1.10 shows both scene-referred and output-referred image preparation for HDR assembly (if we consider the RAW workflow to be a scene-referred one).

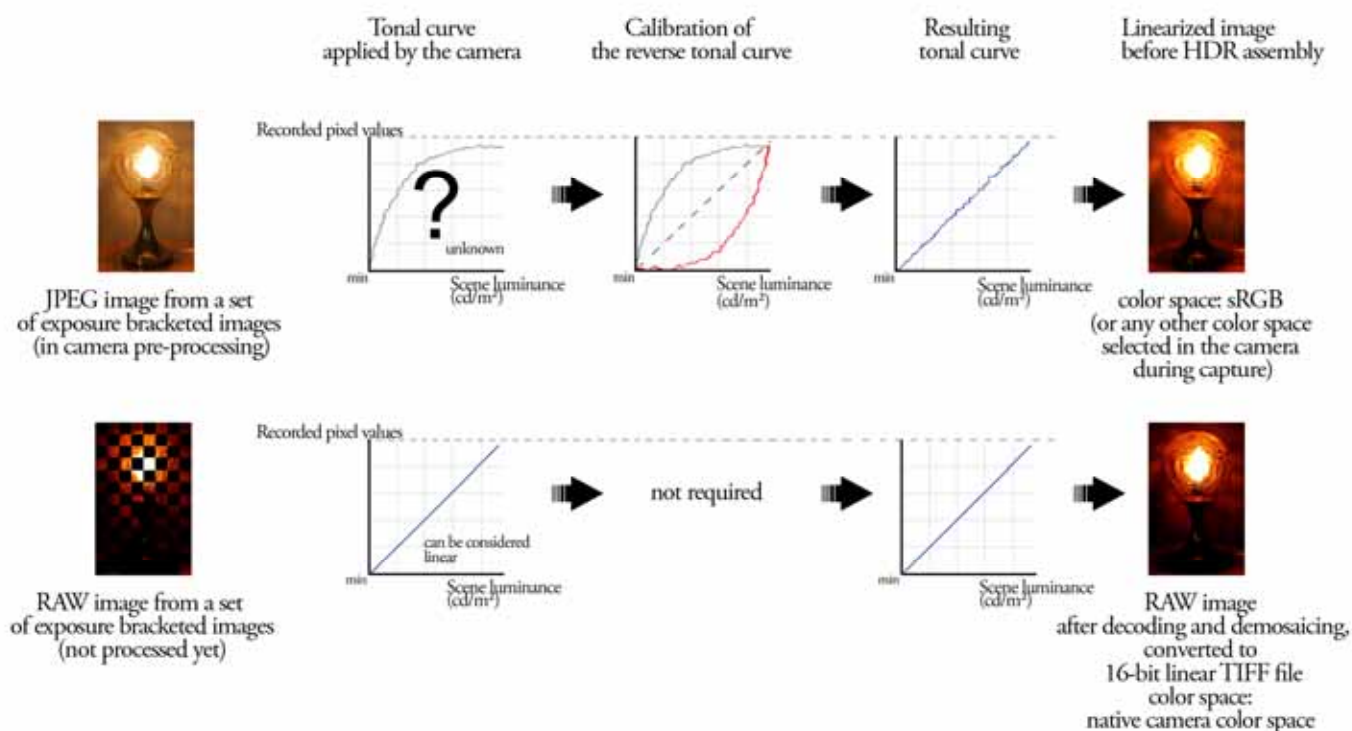


Figure 1.10 A comparison of scene-referred and output-referred image processing prior to HDR assembly

The calibration of the reverse tonal curve from camera pre-processed images always yields an approximation of the true tonal curve applied by the camera and thus "... in most cases, tonal calibration of differently exposed pre-processed outputs of digital cameras is not a reliable process and may lead to inaccurate recovery of the captured tonal values" [22].

Therefore it is best to capture the scene in a camera's RAW file format and subsequently convert the data to linear (preferably 16-bit per channel) images.

1.4.4.1 RAW image data and color spaces

Previous research [21] has shown that the actual color range of a sensor exceeds the widely used color profiles such as sRGB and AdobeRGB. For that the raw camera RGB values were related to the device independent CIE-XYZ color space. "Standard HDR image formats do not impose any constraint on the size of the gamut that can be properly quantized. Our results show that HDR images constructed from characterized unprocessed data can represent at least twice as much visible colors than when constructed from processed outputs with sRGB as color space." [22] Below we provide an example of how the native color space of the Canon® 1Ds Mark II camera relates to the "ProPhoto RGB" color space. Even though the "ProPhoto RGB" space is the largest working color space to encompass the visible color gamut (as defined by the CIE XYZ space), some areas of the camera's native color space are out of the ProPhoto gamut.

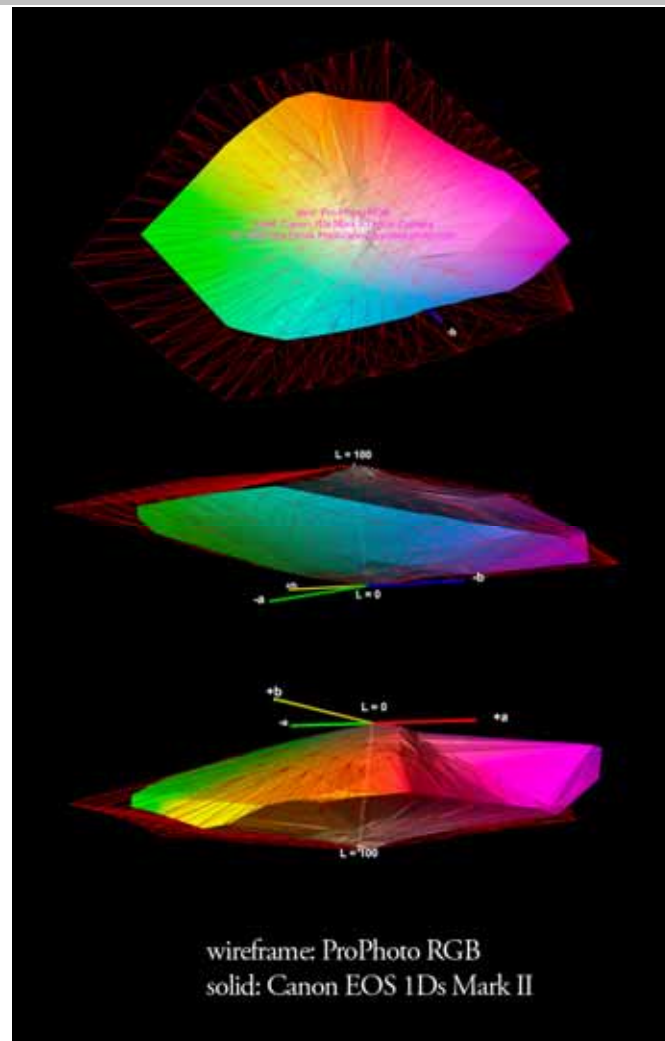


Figure 1.11 Color space comparison in three-dimensional LAB representation
This and other camera space comparisons can be found in [23].

Though tone-mapping issues may arise with HDR images assembled from images with the full gamut of the sensor (values outside a specific working color space like sRGB may get clipped as in [22]), HDR images themselves do not have a specific color space assigned to them (they are not limited by a specific gamut) and thus for IBL purposes the additional color values of the sensor gamut may be of use.

1.5 File formats

After processing high dynamic range image data it has to be stored for further usage. As HDR images contain much more information (in terms of color values) than conventional 8-bit images, they take significantly more space and an efficient compression method is required for their storage. In this section the three lossless file formats most frequently used by the CGI industry are compared. In contrast to lossy compression techniques lossless compression allows to restore the image after decompression so that it is identical to the original.

One of the main qualitative attributes of the HDR image encoding process is the *relative step size* between two encoded values. “For luminance quantization to be completely invisible, the relative step size must be held under 1%” [24]. Relative step size values above this threshold can result in noticeable banding artifacts. Other attributes to consider are the visible gamut coverage, the dynamic range spanned by the encoding process, the bit-depth and the compression type used. Information provided in Table 1.3 is taken from [24] and [25].

File format	Encoding	Bits per pixel	Dynamic range covered (\log_{10})	Relative Step Size	Compression	Coverage of visible Gamut
HDR	RGBE XYZE	32	76 orders	1%	Run-length	no yes
TIFF	LogLuv24	24	4.8 orders	1.1%	None	yes
	LogLuv32	32	38 orders	0.3 %	Run-length	yes
EXR	Half RGB		10,7 orders	0.1%	Wavelet, ZIP	yes

Table 1.3 File formats and their specifications

Further information on each file format can be found in [24]. For HDR images created from bracketed image sequences both precision and dynamic range coverage of all three file types and their encodings are sufficient. In terms of compatibility, HDR RGBE and EXR (OpenEXR file format) are well supported by many popular software packages.

Chapter 2

Mirrored ball acquisition techniques

In the following chapter we shall examine the techniques required to properly capture and process mirrored ball environmental images. In the following discussion the terms *mirrored ball*, *mirrored sphere* and *light probe* are used synonymously.

“Capturing the incident illumination at a point in space requires taking an image with two properties. First, it must see in all directions, in that light coming from anywhere can affect the appearance of an object. Second, it must capture the full dynamic range of the light within the scene, from the brightest concentrated light sources to the dimmer but larger areas of indirect illumination from other surfaces in the scene. In many cases, the standard HDR photography techniques [...] satisfy this second requirement. Thus, the remaining challenge is to acquire images that see in all directions, a process known as panoramic (or omnidirectional) photography.” [26] The first omnidirectional image acquisition technique we shall discuss is the mirrored ball image capture.

The chapter concludes with a set of measurements taken from several randomly selected mirrored spheres to determine their reflective attributes.

2.1 Spherical coordinates and ideal spherical mapping equations

Any environment which may be a subject of HDR acquisition can be represented in spherical coordinates. In such a coordinate system, information on individual points of the environment is provided by two angles φ, θ and the radius of the environment r . Surface points on a mirrored ball can thus be described in this manner. During acquisition, 3D coordinates are projected onto a 2D image sensor plane, therefore if the mirrored ball is assumed to be an ideal sphere the projection of the surrounding environment onto the ball is given by:

$$r = \frac{\sin\left(\frac{1}{2} \arccos(-D_z)\right)}{2\sqrt{D_x^2 + D_y^2}} \quad (\text{Eq. 2.1-1})$$

$$(u, v) = \left(\frac{1}{2} + rD_x, \frac{1}{2} - rD_x\right) \quad (\text{Eq. 2.1-2})$$

with D_x, D_y, D_z being the position of a point in the environment and $u, v \in [0, 1]$ being the coordinates of its corresponding circular 2D projection (as provided in [27]). Conversely, image coordinates of a photographed mirrored ball can be transformed back to the 3D representation by:

$$r = \sqrt{(2u-1)^2 + (2v-1)^2} \quad (\text{Eq. 2.1-3})$$

$$(\theta, \varphi) = (a \tan 2(2u-1, -2v+1), 2 \arcsin(r)) \quad (\text{Eq. 2.1-4})$$

$$(D_x, D_y, D_z) = (\sin \varphi \cos \theta, \sin \varphi \sin \theta, -\cos \varphi) \quad (\text{Eq. 2.1-5})$$

In practice, real-world spheres can be sufficiently approximated by the above equations, provided that the distance from a sphere to the camera is kept large and that the diameter of the sphere is small compared to the environment it reflects.

2.2 Spherical reflections

Frequently it is assumed that a mirrored sphere only reflects half of the environment. However, this assumption is not correct due to the law of reflection: the angle of incidence equals the angle of reflection. This applies to specular (i.e. mirror-like) reflections, which are produced by every mirrored sphere of a sufficient polishing grade. Thus, if a light ray strikes a specular surface at a given point, it is reflected on the surface normal (perpendicular to the point of incidence) at an angle equal to the one at which it


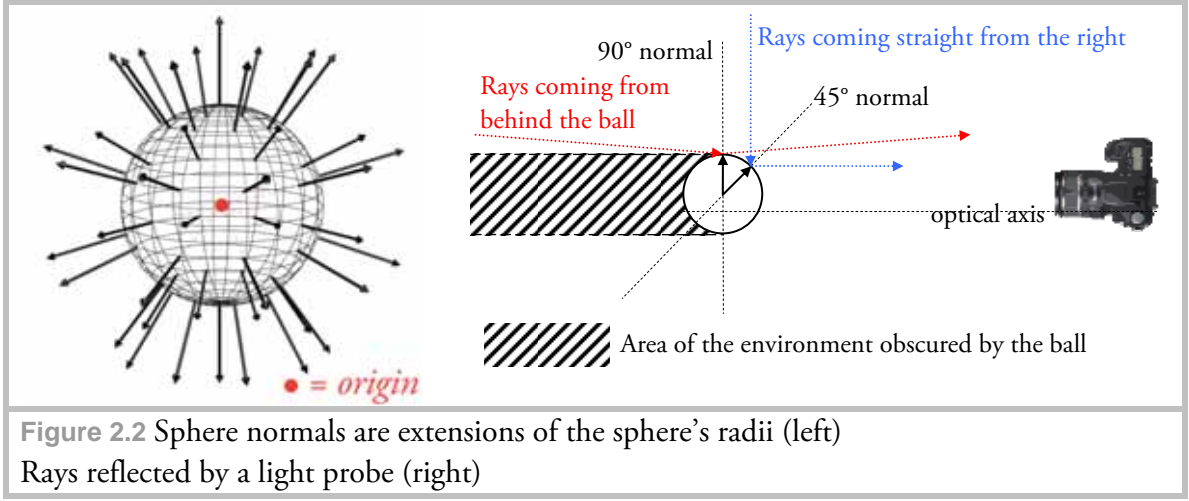


Figure 2.1 Spherical coordinates

reached the surface. Since surface normals of an ideal sphere correspond to radii vectors of the sphere, light rays arriving at the sphere from all directions of the environment can be constructed as shown in Figure 2.2.



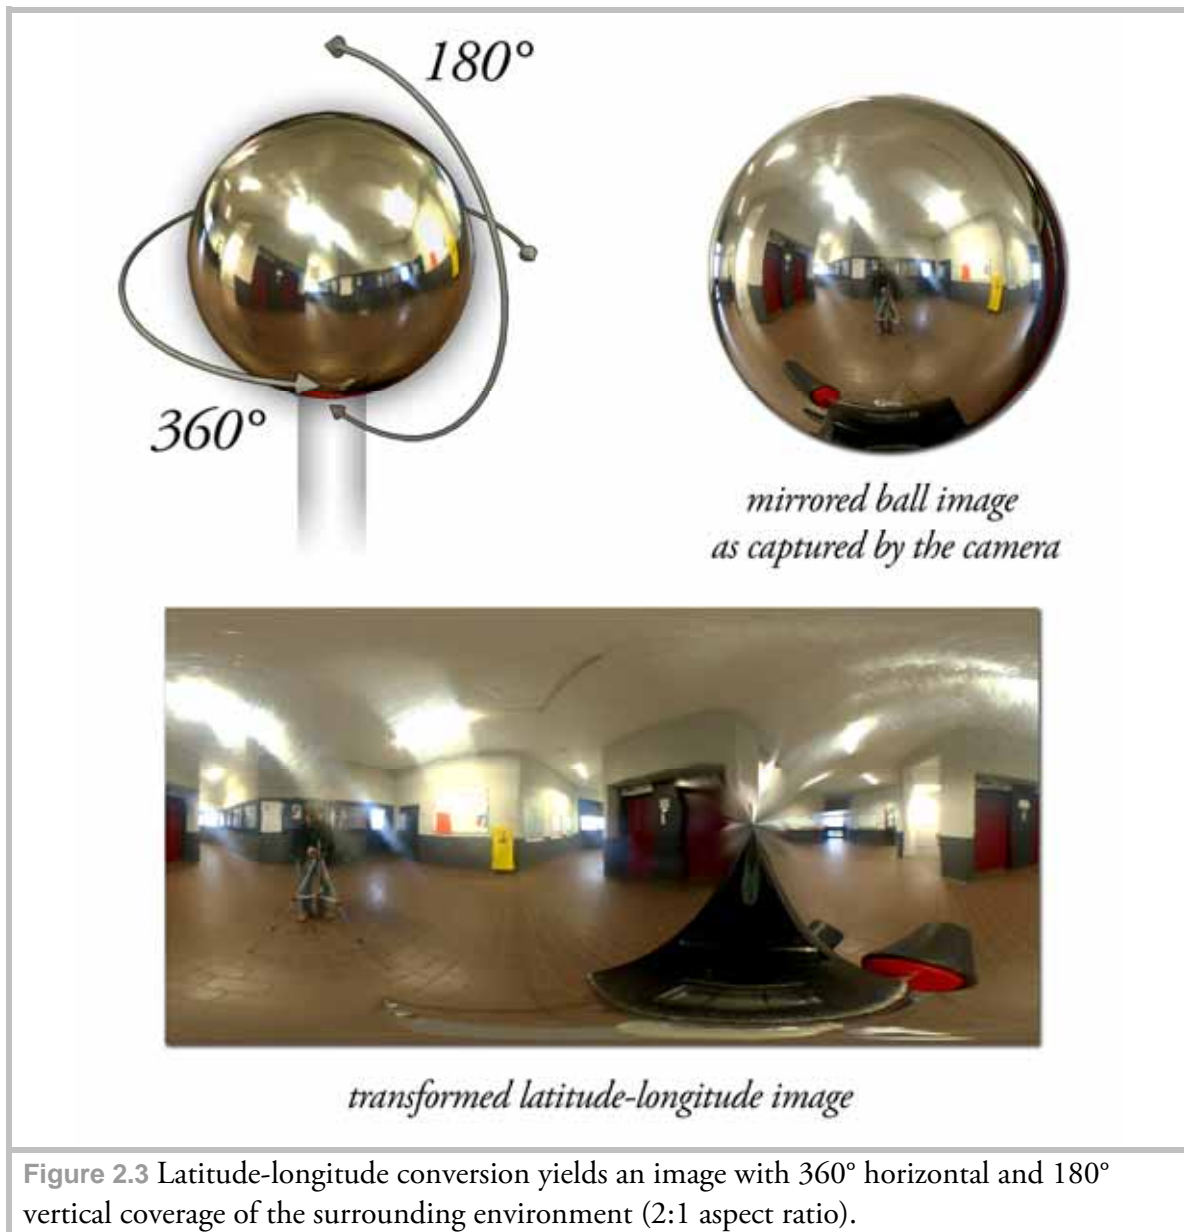
As demonstrated in the Figure 2.2 rays arriving at the outer rim of the sphere (from the camera's point of view) originate from the opposite side of the environment (i.e. almost behind the sphere). Unfortunately the image regions of the outer rim are severely distorted and are thus insufficient for derivation of high quality reflection images. Due to such distortions as well as poor editing capabilities mirrored ball mapping is ill-suited for storage of HDR environment maps and another mapping type is widely used instead.

2.3 Latitude-Longitude mapping equations

Latitude-Longitude mapping (also known as *equiangular* mapping) allows better visual representation of the captured environment and estimation of its quality. In this kind of coordinate system, the environment is mapped to a 2D image by:

$$(u, v) = \left(1 + \frac{1}{\pi} \operatorname{atan2}(D_x, -D_z), \frac{1}{\pi} \arccos D_y \right) \quad (\text{Eq. 2.3-1})$$

with D_x, D_y, D_z being the position of a point in the environment and $u \in [0, 2], v \in [0, 1]$ - the coordinates of the resulting point in the latitude-longitude image (as provided in [28]).



In Figure 2.3 a panoramic transformation of a mirrored ball image reveals some poorly sampled regions (please note that the equirectangular image has been offset horizontally for better illustration). Particularly the area straight underneath the ball (obscured by the tripod head) and the photographer's reflection cover large portions of the image. Thus lighting and reflection information in these areas have to be reconstructed from additional views of the same ball as described in section 2.4.3.

2.4 Before image capture – things to account for

There are several important aspects of mirrored ball photography which should be addressed prior to image acquisition. These are discussed in detail in this section.

2.4.1 Size of the light probe




Figure 2.4 Illustration of the approach described in this section

Often the question arises as to what size the light probe should be and how the resulting reflections relate to the dimensions of the surrounding environment.

The size of the probe affects its physical resolution capabilities. Assuming that objects in the surrounding environment are situated far away from the light probe and that the probe size is small compared to environment (which is a feasible assumption in most cases), the following relation between an infinitesimal surface element dA and its projection dA' on the probe is given:

$$dA = (\sin \theta d\varphi) \left(\frac{R^2}{\cos \alpha} d\theta \right) = \frac{R^2}{\cos \alpha} \sin \theta d\theta d\varphi \quad (\text{Eq. 2.4-1})$$

$$dA' = (\sin \theta d\varphi) \left(\frac{r^2}{\cos \alpha} d\theta \right) = \frac{r^2}{\cos \alpha} \sin \theta d\theta d\varphi$$

where:

dA an infinitesimal surface element in the environment

dA' projection of dA onto an infinitesimal surface element of the lightprobe

R is the radius from dA to the center of the lightprobe (= estimated radius of the spherical environment)

r is the radius of the lightprobe

α is the angle between the normal of the infinitesimal surface element and the direction to the origin of the coordinate system

Figure 2.4 outlines the described quantities. Please note that in the following calculations, infinitesimal surface elements are approximated by planes with no respect to surface curvature. Also, the surface elements' normals are assumed to be oriented perpendicular to their projections on the light probe, hence definitions in Eq. 2.4-1 can be simplified:

$$dA = (\sin \theta d\varphi) \left(\frac{R^2}{\cos(0)} d\theta \right) = R^2 \sin \theta d\theta d\varphi, \text{ and similarly for } dA' \text{ (Eq. 2.4-2)}$$

If the size of a particular object and its reflection (that is, projected area) in the light probe is to be determined, the infinitesimal simplification is no longer valid since in real world applications any area has a certain extent. Thus we have to integrate over the specific area:

$$A = R^2 \cdot \int_{\varphi_1}^{\varphi_2} d\varphi \cdot \int_{\theta_1}^{\theta_2} \sin \theta d\theta = R^2 \cdot (\varphi_2 - \varphi_1) \cdot [-\cos \theta] \Big|_{\theta_1}^{\theta_2} = R^2 \cdot (\varphi_2 - \varphi_1) \cdot (\cos \theta_2 - \cos \theta_1),$$

and similarly for dA' (Eq. 2.4-3)

We have seen that according to the intercept theorems the light probe reflects nearly the entire environment in a single hemispherical projection. Distortions are especially high at the edge, so that only half of the reflected environment can be recovered with sufficient quality.

Thus to obtain acceptable image quality several views need to be captured, and based on aforementioned considerations we can assume for the purposes of our discussion that in each view only half of the environment is being reflected.

We can now relate the environment to its projection on the light probe through the following expression:

$$2\pi r^2 = \frac{2\pi R^2}{\chi} \quad \text{where:} \quad \begin{array}{ll} 2\pi R^2 & \text{is the surface of the a hemispherical environment} \\ 2\pi r^2 & \text{is the surface of a hemispherical part of the lightprobe} \\ \chi & \text{scaling coefficient} \end{array} \quad (\text{Eq. 2.4-4})$$

Eq. 2.4-4 can be solved for r as:

$$r^2 = \frac{R^2}{\chi} \Rightarrow r = \sqrt{\frac{R^2}{\chi}} \quad (\text{Eq. 2.4-5})$$

This equation does not include R' , however it can be neglected without loss of approximation quality if the surrounding environment is big compared to the light probe.

If the spot on the light probe facing the camera is considered to be a planar reflection we can estimate the area of the portion of the light probe's surface covered by an object's reflection visually by measuring its length and width with a ruler. Having an estimated distance to the reflected environment and knowing the approximate size of the desired reflection for a known object, we can thus calculate the radius of the light probe.

Alternatively for a given light probe diameter, the estimated distance of the environment and the size of the reflected object, we can estimate the reflection quality.

For example, let us assume that a light probe with the diameter of 100 mm is available and a 20 m distant environment is to be captured. According to Eq. 2.4-5 we can estimate the reflection size of an object whose area is 1 m² as follows:

$$0.05 \text{ m} = \sqrt{\frac{(20 \text{ m})^2}{\alpha}} \Rightarrow \alpha = 159201; 159201 = \frac{1 \text{ m}^2}{dA'} \Rightarrow dA' = 0,0000063 \text{ m}^2 = 6,3 \text{ mm}^2$$

This example also demonstrates the problem of acquiring high quality environment maps with this technique in large environments: the greater the extent of the environment, the bigger the light probe must be.

2.4.2 Distance between the light probe and the camera

The ideal spherical mapping formulas described in section 2.1 require an orthographically projected image of the light probe on the sensor to deliver accurate transformation results.

Figure 2.5 schematically represents the effect of an orthographic projection. In this example a simulated view through a 28 mm lens (on the left) is compared to a hypothetical view through an imaginary 2000 mm lens. As in the hypothetical projection the virtual sphere is projected onto a tangent camera plane (i.e. the sensor is perpendicular to the principal projection axis) almost exclusively by parallel rays, no foreshortening occurs. Since the aim area of the camera is strongly magnified, the distance to the virtual sphere has to be increased to keep the whole sphere within the visible picture.

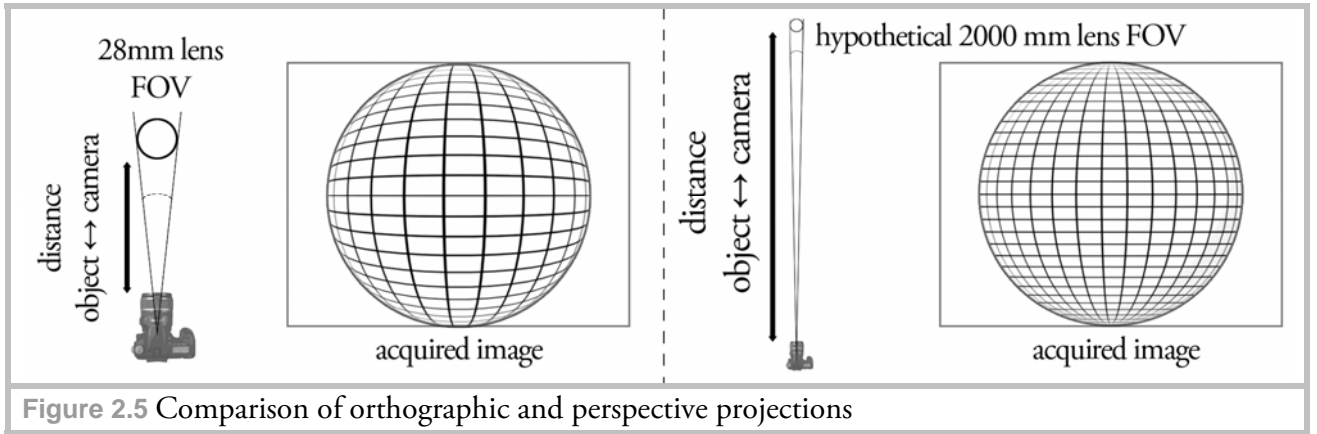


Figure 2.5 Comparison of orthographic and perspective projections

To frame the light probe from a long distance in a real world application a telephoto lens is required. The angle of view, frequently called field of view in photography (FOV), relates the angular extent of the visible environment to the focal length of the camera by:

$$\alpha = 2 \arctan \frac{x}{2f} \quad (\text{Eq. 2.4-6})$$

where α is the FOV angle, x – a dimension of the projection plane (i.e. the sensor) and f the focal length of the lens. As framing the light probe is limited by the vertical sensor size, we are only interested in the vertical FOV (as FOV can be measured vertically, horizontally and diagonally). The following calculations are a simplified model for rays passing only through one lens. These calculations become more complex for zoom-type lenses, but for general understanding the following representation is sufficient.

To properly focus on a distantly placed light probe the camera's focus is set to infinity. Then, following equations hold:

$$\tan(\alpha / 2) = \frac{s_v / 2}{f}$$

$$\Rightarrow \alpha = 2 \arctan \frac{s_v}{2f}$$

(Eq. 2.4-7)




Figure 2.6 Camera attributes involved in FOV calculations

Thus the narrower the field of view (i.e. the greater the focal length), the farther away a light probe can be placed and the better its orthographic projection will be. The distance also depends on the light probe's size: the bigger the light probe, the farther away the camera can be placed. This also results (compare to section 2.4.1) in a smaller reflection of the photographer which simplifies the editing of the resulting HDR image.

2.4.3 How many views to capture

The answer to this question depends on the visible artifacts introduced by various factors such as:

- unevenness of the mirrored ball's surface
- polishing marks and scratches
- objects obscuring important parts of the environment etc.

For a light probe with sufficient geometric (i.e. how close the light probe resembles an ideal sphere) and optical (i.e. how much light is reflected) surface quality capturing two views is sufficient. Although these are usually acquired at a 90° angle to each other, depending on the editing software deviations are allowed.

Adobe® Photoshop®, for instance, has an integrated filter “Offset” which allows for visual




Figure 2.7 Additional views can be taken by rotating the camera around the light probe's centroid

matching of two views in a layered document. Figure 2.7 schematically demonstrates this process.

It is also important to remember that the distance to the light probe r_{camera} must remain constant during capture.

2.4.4 Focusing on the light probe

Similar to the case of the Canon® EF 70-300 mm 4: 5.6 zoom lens which was used for testing, it may prove difficult with many other telephoto / zoom lenses to focus correctly on the light probe [29] In this case close-up filters can be used to bring the whole light probe into focus. For tests in this thesis a BOWER® +1 diopter close-up filter was used.

2.4.5 Optical attributes of the light probe

Apart from the roundness and reflectivity of the light probe the second most important factor are the (micro) scratches and other surface imperfections, such as dents, which can drastically decrease the image quality and require a lot of post-processing of the resulting HDR image. Hence, care has to be taken when selecting a mirrored ball for HDR acquisition. Figure 2.8 shows a scratched light probe and a virtual sphere lit with the HDR image acquired from this probe. Note how a chrome-type virtual material reflects all the surface imperfections of the light probe.



Figure 2.8 Light probe with visible scratches (left)
Threshold filtering the image reveals the scratches around concentrated light sources (middle). Resulting CG image (right)

2.5 During capture



Figure 2.9 Recommended light probe setup with two tripods and exact alignment

In summary following guidelines should be followed for proper HDR acquisition:

- usage of Shutter Priority mode with a fixed Aperture (apertures of F.8 – F.22 normally yield sufficient depth of field);
- maximum manageable distance to the light probe / use of a long telephoto lens (300 mm and above are desirable);
- usage of tripods for stable image acquisition;
- both the light probe and the camera should be on the same optical axis. Usage of a spirit level helps to accurately position and adjust the camera;
- for the best image resolution the light probe should fill the entire frame (i.e. fit into the vertical frame dimension);
- post-adjusting the white balance of an HDR image during the editing step yields much better (in this context: natural-looking) results compared to LDR images [30]. Nevertheless it is useful (and important if images are not captured as RAW data) to white balance the images during the shooting procedure, as described in [31];
- spot metering the scene (either with the integrated camera spot metering system or with an external spot meter) shows the number of exposures to be taken;
- during tests it was estimated that when shooting RAW images an EV step of 2 is sufficient. On the other hand smaller EV steps may help reduce noise artifacts in the final HDR [32];
- very bright light sources are hard to capture as even at the shortest exposure time the sensor pixels become saturated. Particularly the sun's intensity can not be captured with a basic camera and lens setup. Techniques for capturing high intensity light sources are provided in [33];
- moving objects produce ghosting artifacts in the final HDR image. Hence, camera shaking and movement should be avoided, and exposure times greater

than 2 seconds should be used for image capturing (this value has proven to produce acceptable HDR images during testing on a cloudy day).

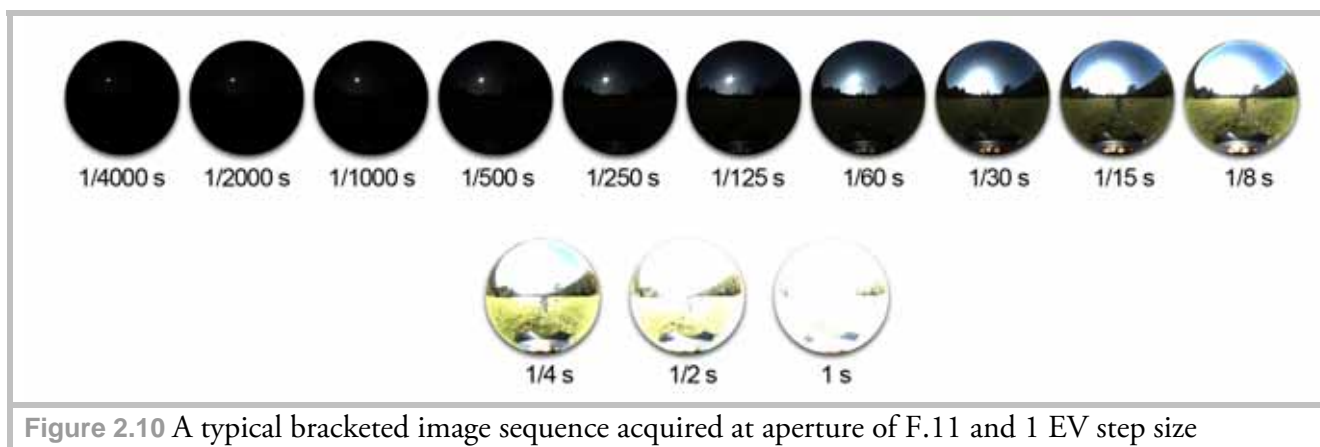


Figure 2.10 A typical bracketed image sequence acquired at aperture of F.11 and 1 EV step size

2.6 After capture

An overview of the processing steps required after the image set acquisition are depicted in Figure 5.4 in Chapter 5. Here we presume a RAW workflow as it is highly recommended.

Different software solutions are available for assembling bracketed image sequences. Since the required post-processing time is critical in a production environment, several solutions were tested, and Photomatix® Pro (<http://www.hdrsoft.com/>) proved particularly useful, as it includes several special features:

- an option to specify a linear-type response curve (as discussed in section 1.4.4);
- a batch processing tool;
- a light probe unwrapping option which produces a latitude-longitude environment map. This – apart from the “angular” map type – is one of the most common mapping types for CGI purposes;
- a user-friendly graphical user interface (GUI).

These features extensively reduced the production time of the final HDR light probe image used for testing in Chapter 5. Settings used for HDR assembly are provided in Appendix C.

2.6.1 Dealing with lens flares

One additional editing step which may be required for all HDR acquisition techniques presented in this thesis is lens flare removal. There are many different factors (quality of the lens and number of lens elements involved, the camera body, the aperture etc.) that may contribute artifacts such as optical diffractions and halos in the final HDR image. Although some automated lens flare removal techniques are available [34], for IBL

applications it is generally sufficient to manually remove the visible artifacts by editing the HDR image in an appropriate software (such as Adobe® Photoshop®).

2.7 Determining the reflective attributes of the light probe

2.7.1 Purpose of the measurements

Reflectivity measurements are not only useful in terms of choosing a suitable mirrored ball, but are a necessity in production when a virtual object's lighting and material attributes are to be matched to a real world photograph (or motion film footage) taken at the same location.

When we capture a scene through a lens the ratio of light transmitted through the lens to the sensor can be considered 100 percent compared to the ratio of light reflected from a mirrored ball. As in production the acquired light probe is eventually used not only for the lighting applications, but also as a reflection map for specularly reflecting surfaces, it is of great importance to calibrate the image. We find the following mention of this matter in [35]:


“It is important to account for the fact that mirrored spheres are generally not optically perfect reflectors. Though the effect is often unnoticed, ball bearings typically reflect only a bit more than half of the light hitting them. In some IBL applications, the lighting is captured using a mirrored sphere, and the background image of the scene is photographed directly. To correct the sphere image so that it is photometrically consistent with the background, we need to measure the reflectivity of the sphere.”

The authors provide a technique to measure relative reflectivity by taking images with a radiometrically calibrated camera. As discussed in Chapter 1 radiometric calibration is easily achieved through a RAW workflow.

The measurement technique provides the relative reflectivity for each individual color channel. Before assembling the HDR image from differently exposed images, for each exposure individual channels can be divided by the measured reflectivity values to correct the images, as if the light probe were 100 percent reflective.

2.7.2 Basic measurement principle

When measuring reflectivity it is desirable to use an adjustable light source of constant intensity as well as a stable rig and a computer generated target (to simplify the isolation process), which will be reflected by the light probe. A simple setup is shown in Figure 2.11.



For the measurements a KAISER® Copy Stand with four 36 Watt OSRAM luminaries (2900 lm, 5400 K color temperature) was used. High intensity light sources help to evaluate the surface of the light probe. The camera angle is not important for the measurements.

As proposed in [35] the averaged pixel values of the photographed target area (which we refer to as “paper patch”) are divided by the averaged pixel values of its reflection (“probe patch”), which yields the relative reflectivity of the light probe.

2.7.3 Tested sphere types



Figure 2.12 Tested spheres (left) and their dimensions (right)

In our experiments, we measured reflectivity values for a variety of sphere types often used for light probe image acquisition:

- stainless steel yard globes: 150 mm, 100 mm;
- Chinese baoding chrome balls: 45 mm, 55 mm (from each set only one ball was measured);
- silver Christmas ball: 90 mm;
- precision steel bearing balls: 75 mm, 65 mm, 60 mm, 55 mm, 45 mm, 35 mm, 28 mm.

Yard globes

Yard globes (also known as gazing balls) are typically used as lawn ornaments and can be inexpensively acquired at garden supply stores. Both low-weight hollow glass and stainless steel variants are available. While glass balls are normally free of scratches they may exhibit double reflections and are fragile. The steel balls on the other hand provide a more robust solution, although in most cases scratches and dents on the surfaces are inevitable.

Chinese baoding chrome balls

Chinese baoding balls (frequently referred to as Chinese meditation balls) are hollow balls of approximately 35-100 mm diameter. They are used in Chinese medicine for injury recovery or as an exercise tool. Some balls have a chrome coating which improves their reflective attributes. However, as no manufacturing standards are available and no concrete manufacturers could be found, two ball sets from different resellers were acquired to determine what quality can be expected in a production situation. Balls from both sets had numerous dents and scratches all over their surface which would require a lot of post processing of the resulting HDR image.

Precision steel bearing balls

Although the bearing balls are produced according to the same standards, different polishing and finishing techniques may be subsequently applied to them. Thus, the 60 mm G100 ball was acquired from a different manufacturer. (Please note that due to budget limitations an equal number of balls from both manufacturers could not be tested. Also, the manufacturers' names can not be disclosed.)

Since ball bearings demand high concentricity the contouring accuracy of all tested balls was sufficient for photographic purposes.

Tabelle 3 — Maß- und Formgenauigkeit, Rauheit

Klasse (Grade)	D_w Nennmaße mm		Grenzab- maße ^e µm	t_{Dws} V_{Dws} µm max.	R_a^f µm max.	V_{DwL}^e µm max.	V_{DwA}^e µm max.	I_G S_T µm	Sortenbereich und Sorteneinteilung ^g µm		
	über	bis									
G3	—	12,7	± 5,32	0,08	0,01	0,13	—	0,5	- 5 bis - 0,5	0	+ 0,5 bis + 5
G5	—	12,7	± 5,63	0,13	0,014	0,25	—	1	- 5 bis - 1	0	+ 1 bis + 5
G10	—	25,4	± 9,75	0,25	0,02	0,5	—	1	- 9 bis - 1	0	+ 1 bis + 9
G16 ^a	—	25,4	± 11,4	0,4	0,025	0,8	—	2	- 10 bis - 2	0	+ 2 bis + 10
G20 ^a	—	38,1	± 11,5	0,5	0,032	1	—	2	- 10 bis - 2	0	+ 2 bis + 10
G28 ^a	—	50,8	± 13,7	0,7	0,05	1,4	—	2	- 12 bis - 2	0	+ 2 bis + 12
G40	—	100	± 19	1	0,06	2	—	4	- 16 bis - 4	0	+ 4 bis + 16
G80 ^b	—	100	± 14	2	0,1	—	4,0	4	- 12 bis - 4	0	+ 4 bis + 12
G100	—	150	± 47,5	2,5	0,1	5	—	10	- 40 bis - 10	0	+ 10 bis + 40
G200	—	150	± 72,5	5	0,15	10	—	10	- 60 bis - 10	0	+ 10 bis + 60
G300 ^a	—	25,4	± 70	10	0,2	—	20	20	- 60 bis - 20	0	+ 20 bis + 60
G300 ^c	25,4	50,8	± 105	15	0,2	—	30	30	- 90 bis - 30	0	+ 30 bis + 90
G300	50,8	75	± 140	20	0,2	—	40	40	- 120 bis - 40	0	+ 40 bis + 120
G500 ^d	—	25,4	± 75	25	—	—	50	50	- 50	0	+ 50
G500	25,4	50,8	± 112,5	25	—	—	75	75	- 75	0	+ 75
G500	50,8	75	± 150	25	—	—	100	100	- 100	0	+ 100
G500	75	100	± 187,5	32	—	—	125	125	- 125	0	+ 125
G500	100	125	± 225	38	—	—	150	150	- 150	0	+ 150
G500	125	150	± 262,5	44	—	—	175	175	- 175	0	+ 175
G600 ^d	alle		± 200	—	—	—	400	—	—	0	—
G700 ^d	alle		± 1 000	—	—	—	2 000	—	—	0	—

^a Nach Vereinbarung mit dem Hersteller können in Ausnahmefällen für die Klassen G16, G20, G28 und G300 die halben Sortenintervall-Werte (I_G) bezogen werden.

^b Nicht in ISO 3290 festgelegt; Klasse entspricht der ehemaligen Klasse IV nach DIN 5401:1978-01.

^c Nicht in ISO 3290 festgelegt; Klasse entspricht der Standardgenauigkeit für ungehärtete nichtrostende Kugeln nach DIN 5401-2:1993-11.

^d Nicht in ISO 3290 festgelegt; Klassen entsprechen den bisherigen Klassen V bis VII nach DIN 5401:1978-01.

^e Werte gelten für den mittleren Kugeldurchmesser D_{wm} .

^f Siehe DIN EN ISO 4288 (cut-off); bei kleinen Kugeln, die nicht mehr nach dieser Norm abgedeckt werden, nach Vereinbarung.

^g Unterteilung in Schritten von I_G .

Table 2.1 “Dimensional and contouring accuracy, surface finish”, DIN 5401
The red circle marks the “roughness” attribute

The roughness attribute was of particular interest during the measurements: we had to determine how (and if at all) this attribute affects the reflectivity values. In the DIN the surface roughness (R_a) is defined as the deviation from the ideal geometric surface [36]. In order to establish the relation between roughness and reflectivity, bearing balls of different grades were tested.

Silver Christmas ball

Christmas balls are in essence smaller versions of glass yard globes and thus have the same advantages and disadvantages. As their surface is very thin they are hard to handle in a

production environment and therefore their usage is not recommended. Nevertheless to give a complete overview of widely used sphere types we decided to include the reflectivity measurement for a Christmas ball as well.

2.7.4 White balancing

Although proper white balancing is not important for the measurement (since pixels in both patches involved in the calculations are affected by the same lighting condition, and thus the relative reflectivity value remains the same), it is useful to calibrate the white balance of the camera by a photographic reference card to have a proper visual representation of the captured scene.

2.7.5 Accounting for differently sized projections of the paper patch on the balls

When measurement images for multiple differently sized balls are to be acquired for comparison, reflections of the paper patch will cover a smaller area on smaller bearings (compare to section 2.4.1). Thus, as the difference in size of the probe patches increases, the measurement uncertainties increase as well. To keep the resulting error small and consistent for mirrored balls




Figure 2.13 A sample image of the modified rig with the 28mm and 35 mm balls in place

of up to 90 mm diameter, a double-sided rig was constructed as shown in Figure 2.13 (The gazing balls were not measured in the same way, as the rig would have required complete redesign).

As the 28 mm ball bearing was the smallest one it was used as the reference for other balls. This ball was placed as close as possible to the target (to maximize the target's reflection) on one side of the rig, and the measured ball symmetrically to it on the other side w.r.t. the target's center. The height of the measured ball was adjusted to visually match the paper patch reflection in the 28 mm bearing. Examination in Photoshop® showed the areas covering the probe patches to have similar size.

2.7.6 Reflectivity measurements and error calculation

During the measurements for each tested sphere ten randomly selected surface points were photographed and their reflectivity values were evaluated.

As for comparison the overall reflectivity of the sphere is of interest (and not that of individual color channels), first the averaged reflectivity was calculated from the evaluated reflectivity value for each individual surface sample:

$$\bar{\rho}_{\text{sample (N)}} = \frac{R_N + G_N + B_N}{3} \quad (\text{Eq. 2.7-1})$$

where: $\bar{\rho}_{\text{sample}}$ is the mean of the sampled reflectivity values
 N sample number
 R_N, G_N, B_N individual color values of the calculated reflectivity

Then, the $\bar{\rho}_{\text{sample}}$ values are averaged (per sphere) which yields the overall reflectivity of the sphere $\bar{\rho}_{\text{sphere}}$:

$$\bar{\rho}_{\text{sphere}} = \frac{\bar{\rho}_{\text{sample (1)}} + \bar{\rho}_{\text{sample (2)}} + \dots + \bar{\rho}_{\text{sample (10)}}}{10} \quad (\text{Eq. 2.7-2})$$

The standard deviation given by the root-mean-square deviation of $\bar{\rho}_{\text{sample}}$ values from the mean $\bar{\rho}_{\text{sphere}}$ shows us how evenly the sphere reflects the environment over the entire surface:

$$\sigma_{\text{sphere}} = \sqrt{\frac{1}{10} \sum_{i=1}^{10} (\rho_{\text{sample (i)}} - \rho_{\text{sphere}})^2} \quad (\text{Eq. 2.7-3})$$

σ_{sphere} can then be given as the \pm value in percent.

2.7.7 RAW Conversion

The RAW file format converter DCRAW has been used extensively throughout the testing since it allows linear, native-sensor-color space conversion to a 16-bit file format (TIFF). As DCRAW is a command-line tool, for measurements and subsequent HDR assembly it is impractical to convert each exposure separately at the OS command prompt level. Instead, a graphical user interface was developed as a front-end to DCRAW to aid the conversion process of file folders. Usage of this DCRAW GUI is described in detail in Appendix A.

2.7.8 Preprocessing of acquired images for calculations

After image datasets have been captured (and passed through DCRAW) each individual image has to be prepared for reflectivity calculation. For that the paper and probe patch areas have to be isolated and written to separate files. Although an algorithm can be developed to search for specific regions in the image and extract the data, it is easier to isolate the regions of interest manually in Adobe® Photoshop® and create a scripted function (called “action” inside Photoshop®), which will then automatically process an entire folder containing all sample images for a sphere. The following section assumes Adobe® Photoshop® is used for image processing.

A sufficient approach is to create a bitmap mask (containing only pure black (0,0,0) and white (1,1,1) values) which will automatically render the unimportant image regions black, leaving only patch areas visible. This is achieved by the “Multiply” blending mode which is described in more detail in Chapter 4. These areas are then saved to new files by the Photoshop® action. Figures 2.14 – 2.16 depict the mask creation steps. Here the image has been cropped for demonstration purposes. However, the techniques described here should be applied to the original (i.e. non-cropped) image to avoid misalignments and manual adjustments during patch extraction. The mask is prepared only once, by filling the areas unaffected by calculations with black color. As the mask contains only two color values the “Threshold...” filter found under “Image → Adjustments” menu is well suited for isolating these regions. The threshold levels for each patch must be set individually. To accomplish this, the prepared, black-filled layer is copied and image regions outside the probe patch area are deleted. The threshold filter is then applied to both layers separately. Additional manual adjustments to the mask may still be required. These are best accomplished by means of the pencil tool (using only pure white and black values). To provide better visual response it is helpful to add an additional adjustment layer on top of the layer stack before applying the threshold filter, as seen in Figure 2.18.



Figure 2.14 A steel bearing sample after RAW processing (Gamma has been adjusted to compensate for linearity of the image to achieve proper printing.).




Figure 2.15 The sample image prepared for mask generation




Figure 2.16 The resulting mask after threshold filtering

An action is generated by recording a full procedure of patch extraction once. Then the steps recorded (as shown in Figure 2.17) can be repeatedly applied to all images on a per-ball basis. After the actions for patch extraction have been created it is useful to create droplets out of them. Droplets are executable actions generated by Photoshop®. They improve productivity as multiple files can be processed simultaneously. The drag-and-drop command of the operating system allows multiple files to be assigned to a new droplet task. Additional information on building droplets can be found in Photoshop® documentation.

Individual stored patches can now be used for calculations.

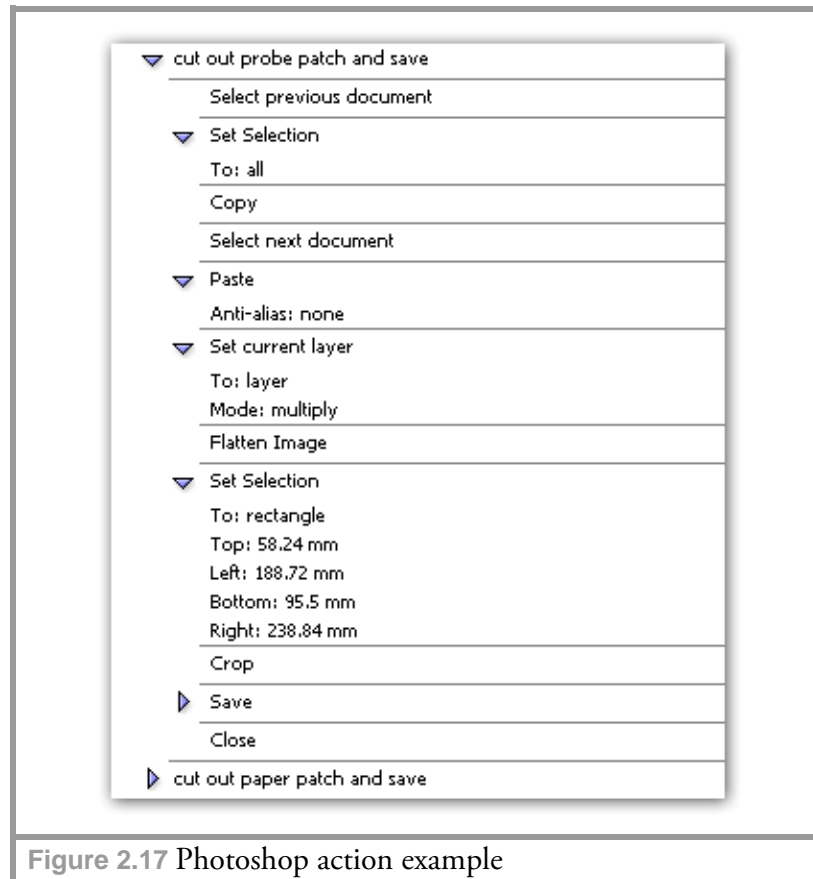


Figure 2.17 Photoshop action example

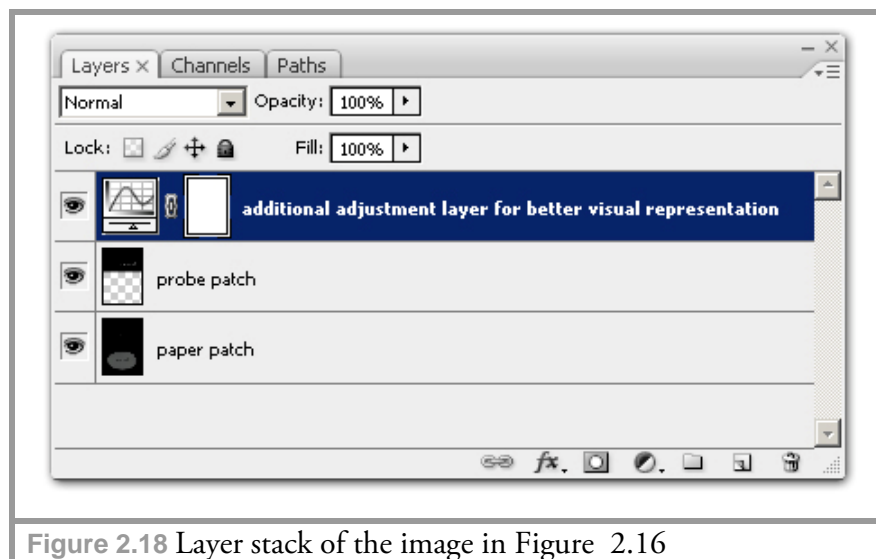




Figure 2.18 Layer stack of the image in Figure 2.16



2.7.9 Calculation results



The extracted patches were processed by the MATLAB® reflectivity calculation tool (as described in Appendix B). The following table is an overview of the calculation results (full version is provided in Appendix B):



Ball type	Averaged Reflectivity
28 mm G bearing ball	55.0 % \pm 0.46%
35 mm G bearing ball	53.7% \pm 0.85%
45 mm G bearing ball	52.6% \pm 0.68%
55 mm G bearing ball	52.3% \pm 0.48%
60 mm G bearing ball	54.2% \pm 0.18%
65 mm G bearing ball	56.4% \pm 0.73%
75 mm G bearing ball	52.4% \pm 1.06%
90 mm Christmas ball	91.5% \pm 1.42%
100mm yard globe	58.5 \pm 1.68%
150mm yard globe	59.0% \pm 0.61%
45 mm Chinese baoding ball	64.6% \pm 0.95%
55 mm Chinese baoding ball	62.1% \pm 1.46%



Furthermore we can provide a reflectivity curve comparison for individual mirrored spheres:















2.7.10 Correcting the bracketed light probe image sequence

During the final step the calculated reflectivity values of a light probe can be applied to the individual images of a bracketed sequence prior to HDR assembly. For this the image color channels are divided by the respective calculated values, which yields an image that would be produced by a hypothetical 100%-reflective sphere. Another MATLAB[®]-based software application has been created for this purpose and is described in Appendix B.

Conclusion

Diffuse lighting information in an environment is easily acquired by capturing any of the measured spheres. As in real world no surfaces are completely diffuse, light probe images have to fulfill additional requirements in terms of reflective quality. Depending on reflectivity attributes of virtual objects in CGI applications, diffuse or specular reflections may be required.

The HDRI acquisition with bearing balls manufactured according to the DIN standard yields the best results in terms of geometric accuracy. On the other hand the 90 mm Christmas ball has the best reflective attributes among the tested probes, but is very fragile. Due to the law of reflection a hemispheric mirror device would be sufficient for full spherical image acquisition, a matter which should be discussed and tested in practice.

The reflectivity of light probe spheres depends on materials used for their production, the sphere type (solid or hollow) and the polishing grade. All of the tested spheres have surface errors which would have required removal prior to CGI application. As none of the tested spheres are meant to be used for photographic purposes, only an average overall quality of the resulting environmental images could be expected. Due to dents and polishing marks all of the tested spheres reflect the environment in the affected areas in a non-specular manner, making such imperfections visible in the resulting image. Thus none of them would have been sufficient in a production environment as any resulting image would require lots of post-work. Although industrial production of bearing balls in Germany runs according to DIN 5401, additional measurements of bearing balls acquired through several different manufacturers would provide a better comparison. When choosing a mirrored sphere as a light probe we need to consider the CGI application it is going to be used in (distance of CGI objects to the camera, animation / still imaging etc.). If high quality light probes are required, special coating (chrome, aluminium) may achieve the desired quality of the acquired images. However, such custom-made light problems are quite expensive, so that fisheye acquisition techniques may provide a good alternative.

Chapter 3

Fisheye lens HDRI techniques

In many situations, high quality environmental images may be required for production. In such cases, HDR acquisition with a fisheye improves productivity and provides the best results in terms of quality.

Fisheye lenses are wide-angle lenses which are capable of projecting up to 180 degrees of the visible environment onto the sensor plane. This lens type is particularly well suited for outdoor HDR image acquisition where rapidly changing lighting conditions demand short capturing times, as its use reduces the number of views of the environment that need to be captured.

The resulting images exhibit variable amount of *barrel distortion* (bending of straight lines) and have to be processed by a distortion correction algorithm before individual images can be combined to a full panoramic image. Fortunately the software presented in this chapter employs the necessary image operations automatically, as do many other similar modern programs.

This chapter gives an overview of the HDRI acquisition process with a fisheye lens.

3.1 Ideal fisheye mapping equations

Several mapping equations are used by lens manufacturers today and here we describe the two most commonly used ones. Both projection types rely on the parameters shown in Figure 3.1.




Figure 3.1 Generic fisheye projection model (simplified schematic)

Most of the modern lenses (such as manufactured by Canon®, Nikon® and Sigma®) are based on *equisolid angle projection*, which looks similar to a mirrored ball projection and is described by the following equation:

$$R = 2f \cdot \sin\left(\frac{\alpha}{2}\right) \quad (\text{Eq. 3.1-1}) \quad [37]$$

The *equidistance projection* was used for production of some older lens types and is given by:

$$R = f \cdot \alpha \cdot \frac{180}{\pi} \quad (\text{Eq. 3.1-2}) \quad [37]$$

If we assign polar coordinates (r, θ) to image points of the resulting projection (on the sensor plane) a transformation into spherical coordinates given in [38] can be applied:

$$\theta_1 = \theta_2 \quad (\text{Eq. 3.1-3})$$

$$\varphi(r) = 1.411269 r - 0.094389 r^3 + 0.25674 r^5 \quad (\text{Eq. 3.1-4}) \quad [39]$$

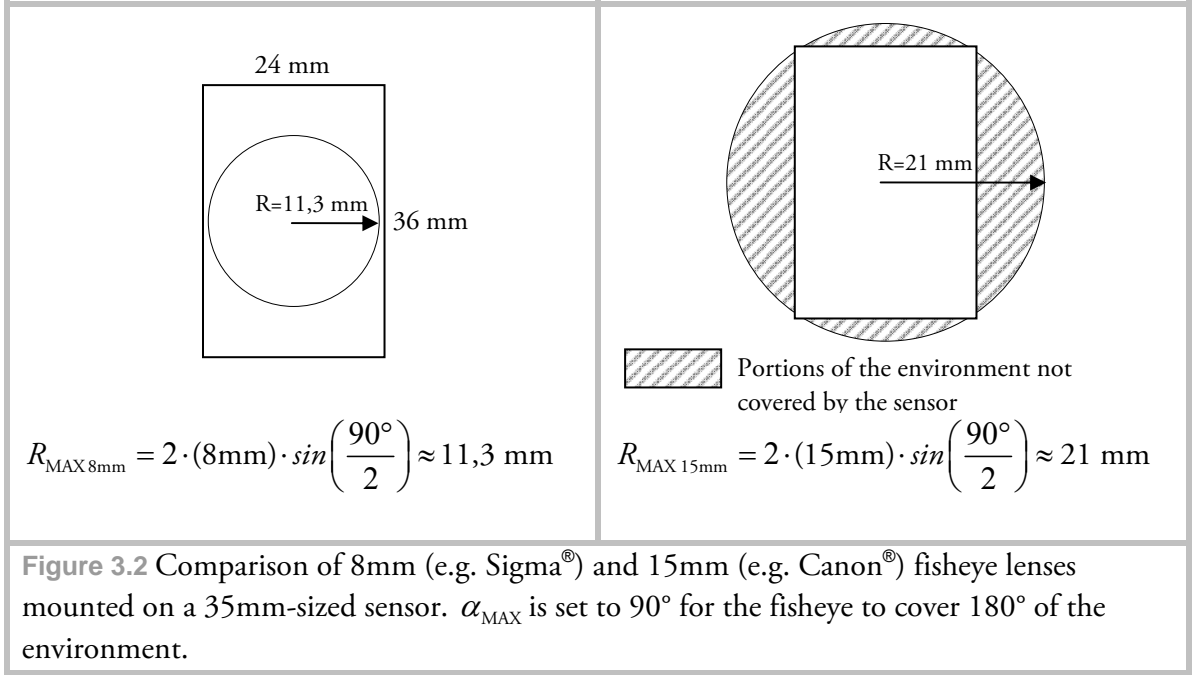
where: θ_1, r – polar coordinates, θ_2, φ – spherical coordinates

Finally the transformation to the canonical Cartesian (x, y, z) coordinate system can be performed as described in section 2.1 (assuming a unit radius spherical environment).

3.2 Fisheye lens type and practical FOV calculations

In practice we are interested in the FOV of the lens as it determines the amount of views to be captured. The fisheye lens type – *circular* or *full-frame* – is directly related to the FOV.

While a circular lens inscribes the projection into the dimensions of the sensor, the full-frame lens circumscribes it. Depending on the sensor size and the focal length of the lens the type of the lens can be determined by applying the Eq.3.1-1. An example is given in the figure below.



As we can see in the figure the full-frame lens projection measures a FOV of 180° only across the frame's diagonal (FOV_D). Thus both horizontal FOV_H and vertical FOV_V calculations become important, depending on the camera setup (portrait or landscape) used to capture the environment.

For the calculations we assume the equisolid angle projection to be valid for most modern lenses:

$$FOV_{\text{equisolid}} = 4 \cdot \arcsin \frac{\text{frame size}}{(\text{focal length} \cdot 4)} \quad (\text{Eq. 3.2-1}) \quad [40]$$

Where *frame size* is the horizontal, vertical or diagonal frame dimension respectively. This ideal formula mostly produces results which deviate somewhat from the manufacturer's specification. For the Canon® EF 15 mm 1:2.8 lens the calculated FOV_D is $184,58^\circ$, while Canon states 180° in their specifications [41]. Furthermore please note that this formula does not hold for circular lenses.

The remaining calculations yield: $FOV_H = 147,48^\circ$, $FOV_V = 94,31^\circ$. Thus in portrait mode $\frac{360^\circ}{94,31^\circ} \approx 4$ and in landscape mode $\frac{360^\circ}{147,48^\circ} \approx 3$ views would be required to

acquire a panoramic image of the entire environment (apart from the additional zenith and nadir views which are needed to cover the poles). It is recommended to capture the environment in portrait camera setup, as in landscape setup additional image rows above and below the horizon level might be required. As panoramic assembly software requires 30-50% overlap for each pair of images one should capture at least two additional images.




Figure 3.3 Required overlap of neighboring images (simplified scheme)

Additionally, the cropping factor of the sensor must be taken into consideration. If for instance an 8 mm circular lens is used in conjunction with an APS-C sensor (22.7 x 15.1mm) its FOV_v gets cropped, as shown in Figure 3.4.




Figure 3.4 Choosing the correct lens for a given sensor size is critical

Therefore apart from lenses designed for use on full frame 35mm cameras manufacturers also provide lenses adapted to smaller sensor sizes, such as the Sigma® 4,5mm F2,8 EX DC circular lens.

3.3 Image acquisition equipment

As fisheye lenses are mainly available for DSLR-type cameras only, a camera of this type is required. Typical lens focal lengths range between 16 and 8 mm. For testing the Canon® EF 15 mm 1:2.8 lens was used.

Furthermore the acquisition process requires several views of the environment to be taken at different angles

while the camera is being rotated around its nodal point (which is normally situated inside the lens). Therefore it is insufficient to rotate the camera around the tripod's vertical axis as this would introduce parallax - a special panoramic head is required. Panoramic heads are available either with an adjustable nodal point, or with the nodal point fixed for a particular lens (<http://www.novoflex.com/english/html/products.htm>, <http://www.agnos.com/>).

One can also construct the head as it was done for this thesis. The nodal point of a lens can be estimated through the techniques described in [42]. Automated exposure bracketing (AEB) can help shorten the capture time. The Canon® EOS 1Ds Mark II camera is capable of capturing up to seven exposures at a time, a feature found only in high end professional devices. If a software development kit (SDK) is available for a particular camera type, a program can be created to remotely control the capture process. Both camera and the tripod should be horizontally and vertically aligned to the plane they are positioned on by several spirit levels (on tripod and camera).

A high speed storage device (either a memory card or IEEE 1394 storage device attached to the camera) is recommended, especially if the AEB capture records more than three images per shutter release. A Transcend® Compact Flash 8GB 133x memory card was used in our experiments for acquisition of both mirrored ball and fisheye image sequences.



Figure 3.5 Custom made panoramic camera head constructed during this thesis

3.4 Post-Processing Equipment

During post-processing, one particular problem lies in image resolution and size (i.e. space in megabytes [MB] an image takes for storage). Canon® EOS 1Ds Mark II captures images at 17 megapixels in the RAW file format (unfortunately, at present time the RAW format can be only used for shooting at full physical sensor resolution). After conversion to the TIFF file format each image can become as large as 100 MB (depending on compression). The converted image sequence first has to be assembled into an HDR

image which in turn may require 25-70 MB to store it (depending on the number of exposures and HDR file format). Thus apart from a powerful computer (dual-core or quad-core system) with at least 2 gigabytes (GB) of RAM required for processing, the respective software used for individual processing steps must be capable of managing such a large amount of data. Additionally, attention has to be paid to the operating system: only 64-bit operating systems support more than 4 GB of RAM. Furthermore the software used for processing must then be 64-bit compliant, as 32-bit applications run on a 64-bit operating system in a compatibility mode which does not allow them to make use of more than 4 GB of RAM.

For processing an Intel-based eight-core personal computer (PC) with 16 GB of RAM with Microsoft® Windows® XP X64 operating system was used. Photomatix® Pro is not available as a 64-bit version yet (and hence was run in 32-bit compatibility mode). Photoshop® CS4 is available in both 32-bit and 64-bit versions, and so is Autopano® Pro. The additional software tools written to aid the conversion and measurement processes do not require a lot of system resources and are completely functional in the compatibility-mode.

3.5 During Capture – Things to account for

Depending on the dynamic range of the scene and the camera capabilities the acquisition process can take up to 20 minutes or more. For instance the Canon® EOS 1Ds Mark II is capable of capturing 19 EVs ranging between 1/8000s and 30s. In order to reduce the time required for capture, it is recommended to spot-meter the scene as described in [43], as it will reveal the actual required EV range (e.g. if according to the metering result the dynamic range is 12 EVs, it won't make sense to capture the full 19 EV range of exposures).

During the tests for this thesis we determined that six views were sufficient to cover the full spherical environment in portrait camera setting with the particular lens and panoramic software (described later) used. Additional views straight above and below the camera might be necessary, although as described in the next section the missing image parts can be filled during the post-processing step. During capturing, the smallest aperture setting (normally up to F 22) should be selected in order to bring the whole environment in focus.

3.6 Post-processing

As the bracketed image sequences were captured in RAW file format the format conversion techniques described in Appendix A apply. Individual views were then assembled to HDR images (see Appendix C). Finally an equirectangular panoramic image was assembled from the HDR images.

Several software packages are available for this purpose and two of them have been tested (Autodesk® Stitcher® and Autopano® Pro). These two applications natively support *stitching* (i.e. panoramic assembly from multiple photographs) of HDR images, whereas other applications require panoramic assembly from LDR images for each individual


exposure prior to HDR assembly. The aforementioned stitching feature can significantly shorten the production time, hence we only tested those software packages that provide it.

Autodesk® Stitcher® proved to be unstable when dealing with high resolution HDR images. In contrast, Autopano® Pro produced reliable results for several high quality HDR environments and offers a complete automated set of tools. Settings required for successful assembly can be found in Appendix D.

3.7 Filling the Nadir / Zenith point

In some situations both automatic and manual image matching which aligns the *nadir* and *zenith* views (the poles of a spherical environment) will not be possible due to insufficient geometric features in the scene. A good example where the automatic algorithms fail is the sky in an outdoor setup: while clouds are moving during capture (or a cloudless sky is present) there are no matching features which could be recognized by the automatic stitching procedure. Manual stitching (also called “forced” stitching) does not produce sufficient results either, as stitching artifacts (like seams and lens vignetting not removed by the stitching algorithms) are likely to show up.

In this situation a manual approach may fulfill the need for a complete spherical HDR environment. Though this technique completely relies on artistic skills of the CGI operator, in most cases the missing portions of the environment are relatively small compared to the overall captured data and are situated at the nadir and zenith points. Hence the obtained results do not negatively impact the overall illumination in a noticeable way and may also contribute to reflections / refractions in a desirable manner. The manual filling technique employed in this thesis is made possible by a sufficient cloning tool with HDR support (i.e. copying parts of an image to fill the missing ones) which has become available Since Adobe® Photoshop® CS3, as well as the Flaming Pear Flexify® 2 Photoshop® plugin (<http://www.flamingpear.com/flexify.html>), which is required for image transformations. To successfully transform a 13346 x 6673 pixel image 4 GB of RAM were needed (tests showed that 2 GB of RAM are insufficient for high resolution images). As the poles of a spherically transformed image are heavily distorted, they can not be edited directly. Therefore a 90 degree rotation around the z-axis (the axis facing the camera, controlled by the “Spin” option in Flexify® 2) is performed prior to editing. After the missing areas have been filled a -90 degree inverse transformation is applied to the edited image.



Conclusion

Fisheye HDR image acquisition can significantly improve the resulting image quality compared to mirrored ball techniques. Although the setup and shooting steps are more elaborate the images produced in such way can be used not only as environment maps, but also as background images.

To become more competitive to automated solutions (such as SpheroCam[®] HDR presented in the next chapter) and produce full panoramic spherical environments in less time a special motorized software-controlled panoramic head could be used. Such a device is available for purchase, called “RODEON VR Head” (<http://www.dr-clauss.de/rodeonvrhead.htm>). The system is capable of automatically rotating the camera and releasing the shutter the appropriate number of times. Additionally, sending the captured data directly to a computer (either by wireless or wired data transfer) speeds up the HDR assembly. Unfortunately due to cost reasons this panoramic head could not be tested in this thesis.

Chapter 4

General compositing techniques

After a light probe image has been acquired and a 3D scene has been successfully lit by the light probe environment, the final output image has to be generated. Although the reflections / refractions may be physically modeled, the artist may decide that additional adjustments are required to achieve the final look, or to meet the requirements of a customer. In many cases these minor adjustments are hard to carry out in the 3d package. Instead a technique called “compositing” is used.

Although CGI imagery can be exported (rendered) straight from a 3D software package, in many cases additional adjustments are required not only to the image as a whole, but to individual elements and portions of the image like highlights or shadows. In this chapter a technique is introduced to break up the image into its individual elements (or components) in order to gain access to post-rendering adjustments. This technique has not only proven to be an important addition to the CGI-only workflow, but it is in fact the only sensible way to combine live action footage (or still photography) with CGI elements to achieve realistic results.

4.1 Compositing

In [44] digital compositing is defined as follows:

“Digital Compositing: the digitally manipulated combination of at least two source images to produce an integrated result.” As far as the main goal of compositing is concerned, the author states: “By far the most difficult part of this digital compositing process is producing the integrated result – an image that doesn’t betray that its creation was owed to multiple source elements. In particular, we are usually attempting to produce [...] images that could have been believably photographed without the use of any post-processing. [...] Even if the elements in the scene are obviously *not* real [...], one must be able to believe that everything in the scene was photographed at the same time, by the same camera.” [44]

4.2 The multipass rendering workflow

Multipass rendering is a way “to render each of (the) multitudes of surface and lighting attributes as separate images, or passes.”[45] The images generated that way can be combined in an external compositing package, yielding the final image. There are many advantages to this technique, among them:

- adjustment to individual objects after the rendering can be accomplished by re-rendering only those objects;
- in some cases individual frames require less time for rendering, since less geometric data is to be processed;
- color adjustments are easily achieved, giving a real-time preview during compositing;
- computationally intensive post-processing work (like motion blur and depth-of-field effects) can be quickly handled during compositing.

4.3 Mathematical background

In order to understand how a multipass rendition is composited into the final image we shall first have to examine several relevant mathematical operators which provide a certain look to a *pass* when it is being combined with other passes (the term *pass* is defined in section 4.4).

In general the bit depth of the resulting image defines how the operators affect the input values. Currently many media productions are based on either 8-or 16-bit integer workflows. Hence the maximum pixel values for each channel (RGBA) are limited to 255 and 65535 respectively. The results of some of the operators can often exceed these maximum pixel values. This causes the overflowing pixel values to be clamped to the maximum allowed value, an effect known as saturation which results in a clamped image. In contrast, 32-bit floating point formats support pixel values with a much larger range, so the resulting values do not get clamped [46].

All of the image arithmetics are applied on the per-pixel basis to the input images. In matrix representation this can be expressed as:

$$iC(m,n) = iA(m,n)(operator)iB(m,n)$$

where iC is the resulting image, iA and iB are the images the operator is to be applied to, and m,n – the pixel to which the operator is applied (this notation is also used in all subsequent descriptions). For all the described operators that follow the order of iA and iB can be swapped.

4.3.1 The Add operator

The Add operator takes as input two images and produces a third image in which each pixel value is the sum of the values of the corresponding pixels from each of the two input images. A common variant of the operator simply allows a specified constant to be added to every pixel [46]. The output image is simply given by:

$$iC = iA + iB \text{ (Eq. 4.3-1)}$$

Figure 4.1 shows an example of the Add operator in action. Zero black values act as transparency in the image added.



Figure 4.1 The Add operator (from left to right): iA , iB and the resulting image iC

4.3.2 The Screen operator

The Screen operator equation is given by:


$$iC = 1 - ((1 - iA) \times (1 - iB)) \text{ (Eq. 4.3-2)}$$

The Screen operator is best suitable in situations where bright areas of an image are to be composited over a background image [47]. A result of the Screen operator in action can be observed in Figure 4.2.



Figure 4.2 Combining images via Screen operator (from left to right): iA , iB and the resulting image iC

For suitably chosen input pixel values of two identical gradient images, a plot of the Screen operator function is given in Figure 4.3 (values used depict only one color channel). We recognize that the operator affects the overall brightness of the resulting image, universally lightening it. Thus the screen image has to have a zero black background in order for the transparency to work. The main difference to the add operator is that the Screen operator “[...] is less likely to reach pure white and suffer from clipping problems in bright areas.” [48] Furthermore it allows a softer blend between the input images.



4.3.3 The Multiply operator

This operator mixes two images together by multiplying their pixel values. The output image is given by:


$$iC = iA \times iB \quad (\text{Eq. 4.3-3})$$



Figure 4.4 The Multiply operator (from left to right): iA , iB and the resulting image iC

Since the pixel values get multiplied, the portions of the image that are to be left unaffected have to be 100 % white.

Plotting the Multiply operator function for two identical gradient images yields the result in Figure 4.5. By looking at the plot we recognize that the Screen operator can roughly be described as the inversed multiply function. While multiplying images generally produces a darker result, the screen operator does the opposite.



4.3.4 The Difference operator

As described in [49] the Difference operator “subtracts the pixel value of (one) layer from that of the (other) layer and then takes the absolute value of the result. The mode is commutative; the order of the two layers doesn't matter. Because the order of the layers doesn't matter, Difference operator examines the numerical values for color in both the base and blend images, and subtracts the lower number (darker color) from the higher number (brighter color). The result color has a numerical value which is the difference between the two values.” The output image is given by:

$$iC = |iA - iB| \text{ (Eq. 4.3-4)}$$



Figure 4.5 The Difference operator (from left to right): iA , iB and the resulting image iC

4.4 Common rendering passes

In the following section several commonly used rendering passes are described, along with their influence on the final result. As opposed to layer rendering where different objects in the scene are rendered separately, rendering in passes addresses different shading components of an object. “The term component refers to one element of a shading model equation.” [50] An example of such equation is given in Eq. 4.4-1 (as described by Bui Tuong Phong in [51]).

$$S_p = C_p [\cos(i)(1 - d) + d] + W(i) [\cos(s)]^n \quad (\text{Eq. 4.4-1})$$

where:

S_p is the resulting shading at point P on an object

C_p is the reflection coefficient of the object at point P for a certain wavelength

i is the incident angle

d is the environmental diffuse reflection coefficient

$W(i)$ is a function which gives the ratio of the specular reflected light and the incident light as a function of the incident angle i

s is the angle between the direction of the reflected light and the line of sight

n is a power which models the specular reflected light for each material

[51]

In modern 3D applications shading model equations are implemented as *shaders* which possess the attributes that describe individual components of the equations (as seen in Figure 4.6).

As such each shading component can be related to a specific visual element of a rendition. “For instance, the specular component of a shiny surface is the specular highlight.” [50]

Depending on the shaders involved in the rendering process the shading components comprise several shader (also called material) and light attributes. “For example, the diffuse component of a surface is determined by the Diffuse attribute of the assigned material and the Intensity attribute of the light hitting the surface. If the surface is assigned to a Blinn material, the specular component is determined by Eccentricity, Specular Roll Off, and Specular Color attributes, as well as the light’s Intensity attribute.” [50]

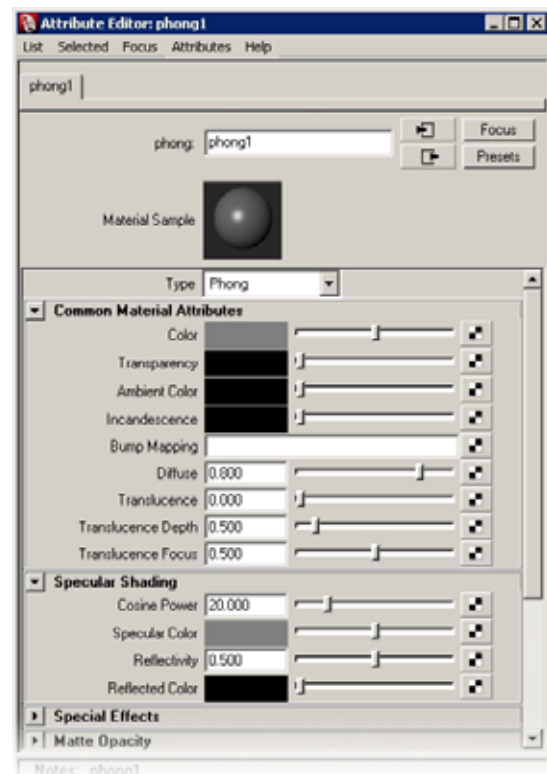


Figure 4.6 The Phong shading model implementation in Autodesk® Maya® 2008

4.4.1 Description of different passes

The following section provides an overview of different render passes.

The overview describes several commonly employed passes, although additional passes (predefined in 3D applications or custom-made) are also widespread in practical applications.

Beauty Pass

“A beauty pass is [...] a complete rendering of [the] scene, complete with all of the attributes such as reflections, highlights, and shadows. [...] A beauty pass is made redundant if [we] render other passes that can be composited together to recreate it.” [48] Nevertheless it is often useful to have the beauty pass rendered for comparison. By doing so the compositor has an overall view of the scene which he / she can reference during the compositing.

Ambient Pass

This pass (so called “color pass” in some applications) produces color and texture maps on the object’s surface but ignores any additional shading information. “An ambient pass shows every object as if it were uniformly lit by ambient light, so objects appear in a consistently flat, uniform tone, [...] There will not be any shading to brighten or darken parts of the surface, as there would be in a diffuse pass. The colors in rendered pixels in an ambient pass will be the exact color taken from the corresponding pixel of a texture

map, or the exact surface color of untextured objects, without any influence from lighting.” [48]

Diffuse Pass

“A diffuse pass is the full-color rendering of [the] subject, including diffuse illumination, color, and texture, but not including reflections, highlights, or shadows, [...] Because a diffuse pass includes the diffuse illumination from lights, surfaces are shaded brighter where they face a light source and darker where they face away from a light source.” [48] As shown in the practical part of this chapter, although a diffuse pass can be rendered precomposed, in many cases it is wise to composite this pass from additional passes, a keylight pass for instance, which in turn provides additional control over the diffuse shading.

Specular Pass

“Specular passes (or highlight passes) isolate the specular highlights from [the] objects [...] Rendering a separate specular pass allows you more creative control over how the highlights are rendered [...] During (the compositing), having specular highlights as a separate pass will allow control over their color and brightness, so that [we] can adjust the highlights to match the rest of [the] composited scene.” [48]

Reflection Pass

“A reflection pass can include self-reflections, reflections of other objects, or reflections of the surrounding environment. Often [we] need to render several reflection passes, especially if [we] want to isolate raytraced reflections on different objects. [...] To set up a reflection pass showing the reflections of the environment onto the [object], the [object] is given a shader that does not show diffuse illumination or specular highlights [...], but is reflective.” [48]

Shadow Pass

“A shadow pass is a rendering that shows the locations of shadows in a scene. In scenes with overlapping shadows, it is important to keep the different shadows separated when rendering shadow passes, so that [we] can control their appearance, color, and softness separately during the composite.” [48] The shadow pass can be separated into objects casting shadows on other objects and self-shadowing of objects.

Key Light Pass

A key light pass provides lighting information from the main light source (e.g. sun) for objects in the scene. It is best used in conjunction with a diffuse (lambertian) shader and composited over the ambient pass to form the diffuse pass.

Global Illumination Pass

“A global illumination pass isolates the indirect light added to [the] scene by global illumination, [...] This way, [we] can adjust the balance of direct and indirect light in [the] rendering.” [48]

Ambient Occlusion Pass

“Ambient occlusion provides a way of adding realism to images by taking into account the influence of occluding objects while determining the color at every given point on a

surface.” [52] Speaking of CGI elements, “[...] it is the process of objects blocking the flow of light from each other, which has an effect of darkening the surface area on an object as it gets closer to another surface. This also affects how a surface reacts to its own geometry such as around folds or creases (self-shadowing).” [52]

“Ambient occlusion shades the scene based on the proximity of nearby objects, without taking into account their brightness, color, or transparency.” [48] In some situations, certain objects, such as luminaries, should be excluded from this pass.

Reflective Occlusion Pass

“Reflective occlusion provides a way to control the amount of reflection color along a surface based on the reflection rays so that it reduces the reflection intensity where surfaces are in close proximity, based on occlusion effects.” [53]

“To leverage the occlusion shader’s output, [we] need to multiply it with the reflectivity value of a given shader so that it affects the overall reflection intensity across the surface. The result would be a reflection shader that scales the amount of reflection across the surface based on reflective occlusion, rather than an overall linear value.” [54]

Z-Depth Pass

“The Z-depth (pass) [...] represents the scene depth based on the camera’s viewing plane. [...] Z-depth is used for depth sorting with per-pixel effects, such as depth fading, 2D fog, or fast depth-of-field post-process effects.” [55]

“Z-depth post-process effects, whether with an output shader or in a compositing application, a high level of accuracy is required for representing a transition over distance as means for avoiding unwanted artifacts, and thus you should always use (and export) 32-bit floating point Z-depth passes.” [55]

The Motion Vector Pass

“The motion vector pass [...] stores motion vectors typically for 2D (post-processed) motion blur effects.” [56] As with Z-depth passes, motion vector passes simplify sensitive adjustments that establish a certain look of motion blur and also reduce its calculation times.

4.5 Compositing example

The following is an example of how multipass renditions of a scene can be composited together using the math operators described in 4.3. It must be mentioned though, that since multipass rendering in many cases is dependent on specific shaders used in a production environment, there is no generic approach for every situation. Hence the example demonstrated uses simplified shading networks and provides only one of many possible solution. As different software packages evolve rapidly only a generic compositing diagram and an example compositing breakdown in Adobe® Photoshop® CS3 is presented here.

Furthermore we have to account for the fact that aside from the math operators described in Chapter 4.3, compositing itself is a rather artistic process, therefore no additional technical background information can be provided here. The compositor decides which look is to be achieved. From here he / she can deduce which operators should be used.

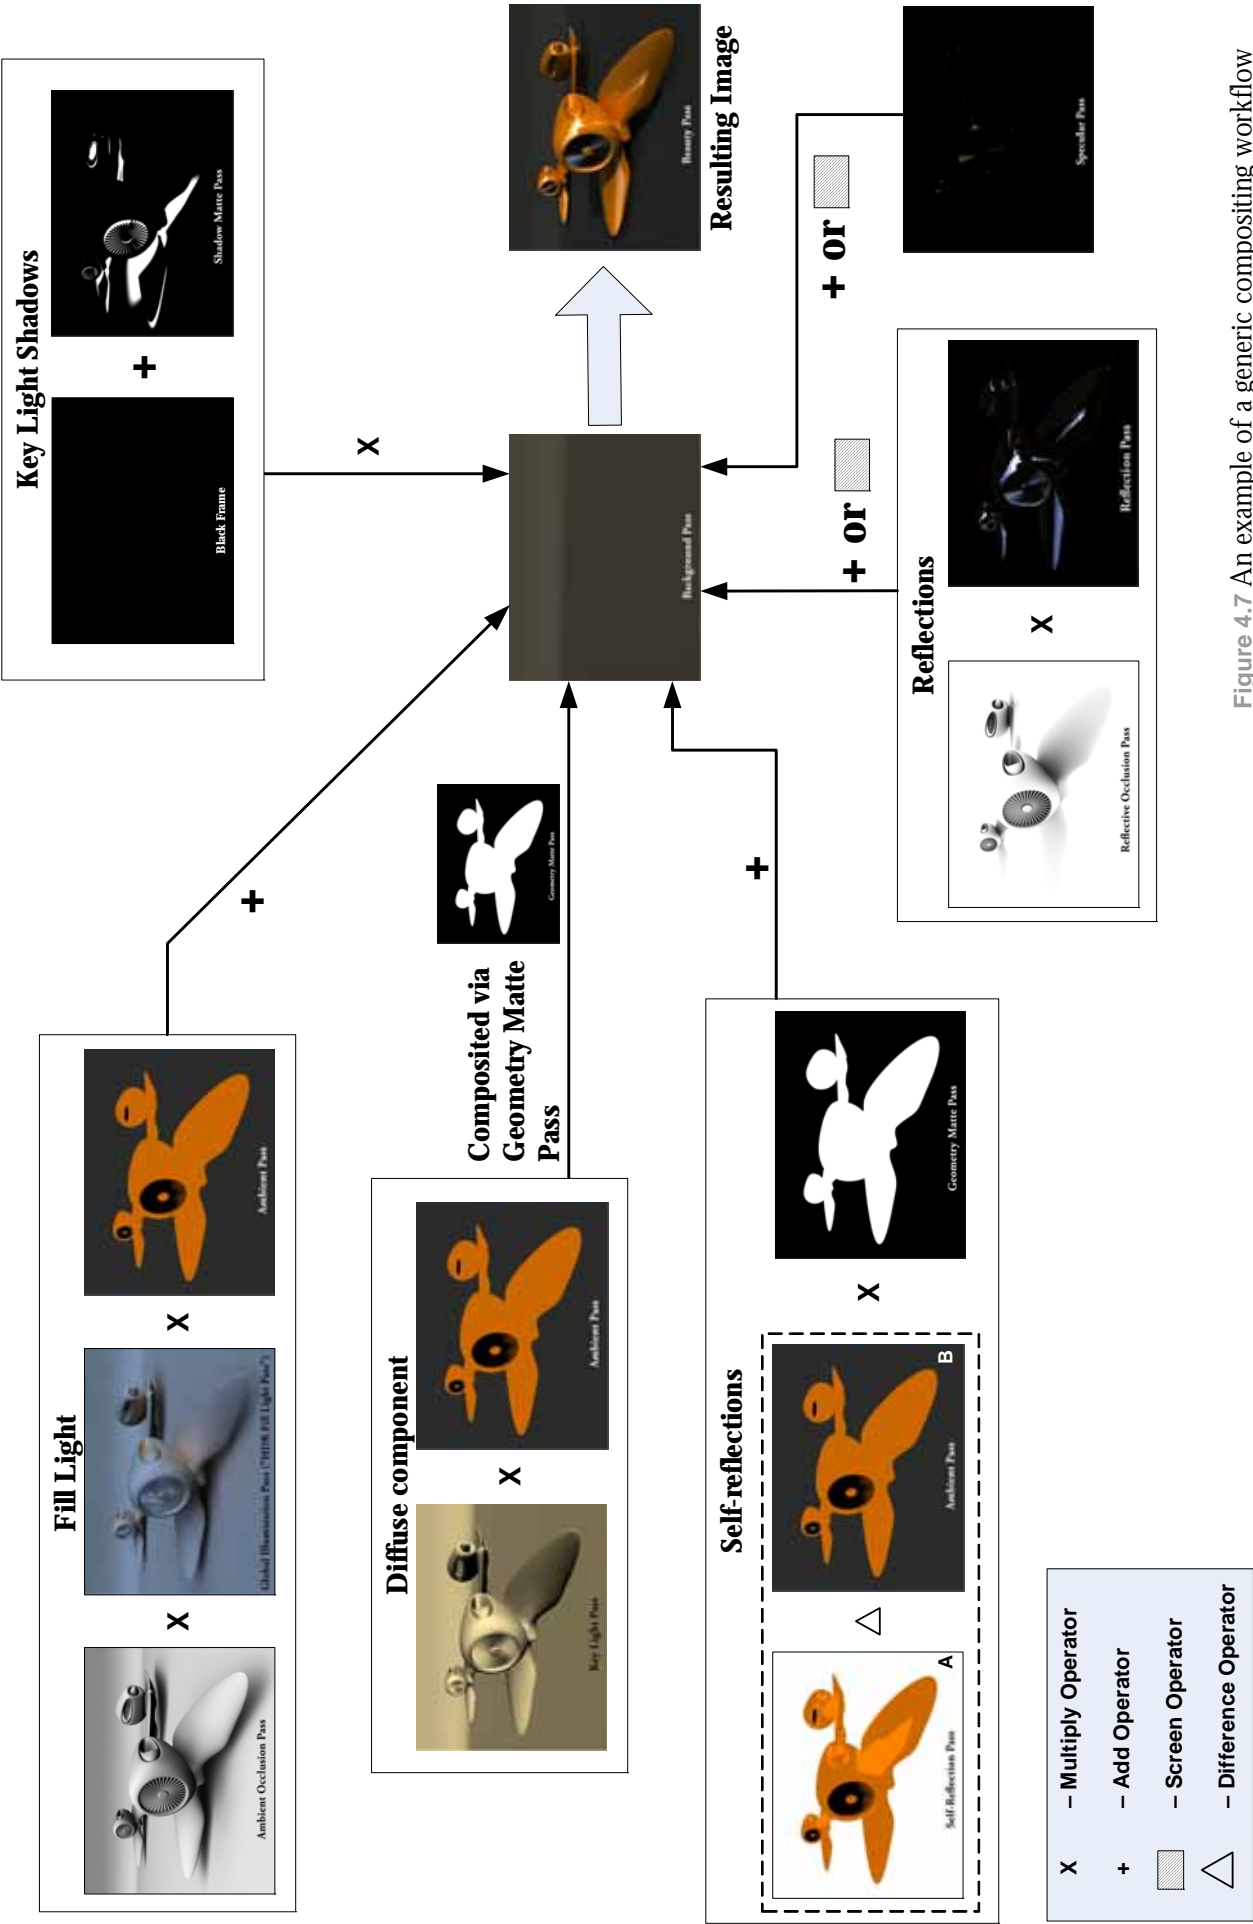
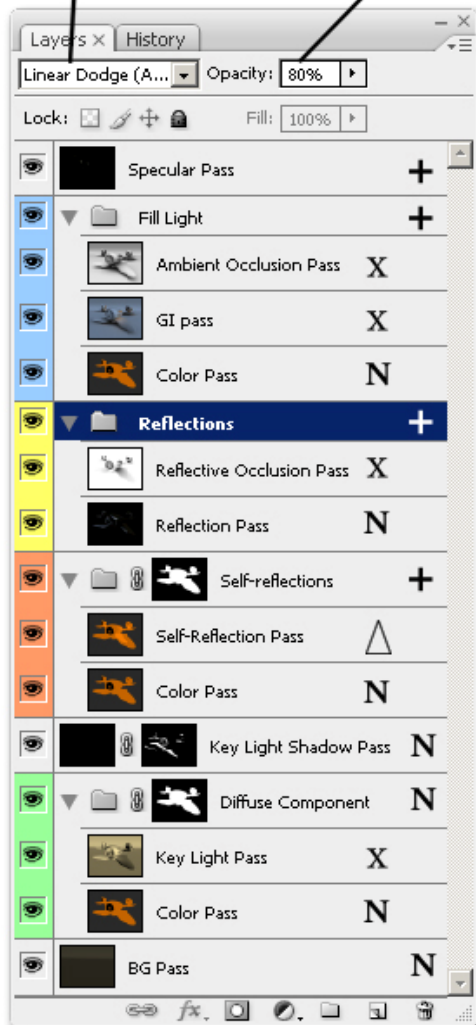


Figure 4.7 An example of a generic compositing workflow

The Blending Modes allow different math operators to be applied to layers, affecting both the selected and the underlying layer

The opacity control allows linear transparency adjustments for individual layers



- + Linear Dodge (Add Operator)
- x Multiply Operator
- N Normal Mode (no operator applied)
- Δ Difference Operator

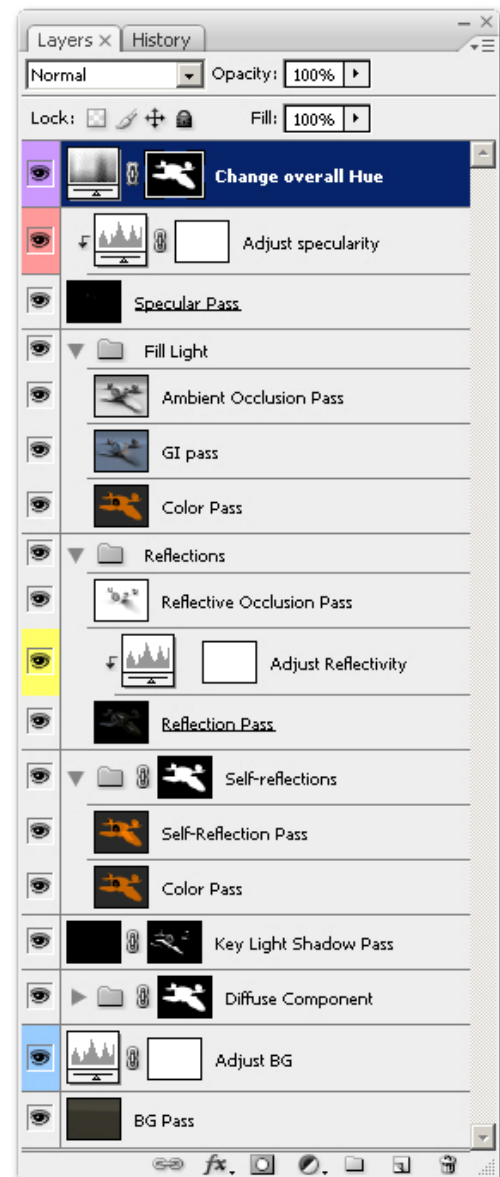


Figure 4.8 An implementation of the compositing scheme (found in Figure 4.7) in Photoshop®

Figure 4.9 Adjustments are easily achieved as post-rendering effects




Figure 4.10 The resulting image after compositing (Figure 4.8 breakdown)



Figure 4.11 The resulting image after compositing (Figure 4.9 breakdown)

In Figure 4.8 individual layers (the render passes) are stacked on top of each other to form the final image. Colored visual aids are provided for layer groups and their respective layers. The grouping concept inside Photoshop® allows nested compositing techniques. Multiple layers are combined inside a “container” (i.e. they are precomposed) which in turn can be composited with underlying passes. This allows a one-to-one implementation of the concept provided in Figure 4.7. Furthermore the masking concept is utilized to separate areas of the scene during compositing. The multiplication of the “Geometry Matte Pass” is essentially done by applying a mask to individual layers or groups.

Figure 4.9 shows how easily adjustments of the resulting image and of each individual layer can be achieved throughout the compositing process (changes are indicated by black arrows). For instance, changing the base color of the object would have normally required the whole image to be re-rendered. Here on the other hand, it becomes a simple matter of an HSV-operator (hue / saturation / value) to be stacked on top of a layer to modify the layers color attributes.

Other compositing packages like Apple® Shake®, The Foundry® Nuke®, Autodesk® Toxic® etc. may provide additional compositing features not discussed here.

Conclusion

This chapter introduced the compositing techniques. As nowadays the technology evolves rapidly new techniques may appear, which provide better render passes in less time. By the time of this writing Autodesk® Maya® 2009 had just been released. This version includes enhanced render layer management and provides many predefined render passes which had to be set up manually in previous versions. Furthermore better integration between 3D and compositing applications can be expected. Autodesk® Maya® 2009 provides superior integration with Autodesk® Toxic®, allowing the render passes to be linked together by math operators already in Maya®, passing a combined render pass stack to Toxic®. Hence artists can be more productive, concentrating on the creative process.

Chapter 5

Comparison of the different HDRI acquisition techniques

In this final chapter we provide a comparison of the introduced HDR image acquisition techniques.

The main determining factor in HDR environmental imaging is the dynamic range a capturing device (or system) can cover. Although several HDR imaging sensors are available on the market [57], their resolution is insufficient for high quality digital capture, and production of high resolution HDR sensors is still expensive. The next section compares dynamic ranges covered by the fisheye technique described in Chapter 3 and by an integrated professional solution.

As the techniques differ significantly in terms of costs, processing time and image quality the differences in virtual scene lighting and reflections produced by the resulting HDR are of interest. The company 809cgi GbR who supported this thesis provided a 3D model of an Audi R8 car and the photographic location where a comparison could be made. The results are presented in this chapter.

The fisheye technique may become competitive to the professional solution through use of advanced technology (such as automatic acquisition software). The proposal for such an approach is made at the end of the chapter.

5.1 The SpheroCam HDR

Several panoramic capture solutions are available on the market today. While most of them are primarily aimed at photographic usage the SpheroCam[®] HDR is of particular interest for the CGI industry. At the time of this writing, it is the only available fully automated solution for acquisition of full spherical HDR environments. The following technical specification is taken directly from the SPHERON-VR AG website [58]:

- captures in vertical scanlines, while rotating about the vertical axis (one-shot 360°x180° HDR scanning, no stitching, no bracketing exposures);
- scan time: 20 seconds up to several minutes, depending on exposure time, chosen dynamic range and resolution;
- up to 50 megapixel (5.300 x 10.600 pixel) spherical and up to 100.000 pixels horizontal panoramic resolution;
- dynamic range: up to 26 f-stops;
- exposure times: 1/100.000 up to ¼ seconds per scanline (1/6 f-stops in between standard graduation);
- 96-bit technology: Radiance (.hdr) and OpenEXR (.exr);
- lens: Nikon[®] Nikkor AF-D 16 mm (individually) calibrated fisheye lens for spherical images, up to 200 mm lenses for panoramas.



Figure 5.1 SpheroCam[®] HDR, with permission from 809cgi GbR

The SpheroCam[®] HDR is delivered with proprietary software, which is shown in Figure 5.2.

Another solution, the Panoscan[®] MK-3 system offered by Panoscan[®] Inc., captures only up to 11 EVs of dynamic range but claims to deliver better image quality [59]. The 809cgi GbR company kindly provided their SpheroCam[®] HDR for comparison purposes in this thesis.




Figure 5.2 Software interface of the SpheroCam[®] HDR

5.2 Comparing the techniques

As the overall capturing and processing time of images / image sets acquired by each technique depends on the *operator's* skills and on computer performance, no precise time comparison can be given here. We do, however, provide an overview for each acquisition technique, thus outlining the complexity of these, as shown in Figure 5.4.

For the test described in this chapter all three techniques (mirrored ball, fisheye and SpheroCam[®]) were applied to acquire three equirectangular HDR environmental images of the same location (as shown in Figure 5.5). Average time requirements for each technique are given as reference values in Table 5.1. Please note that the reference values for both fisheye and mirrored ball techniques are provided w.r.t the technical equipment used in this thesis:

- Fisheye: Canon[®] EOS 1Ds Mark II (17 megapixel), Canon[®] EF 15 mm 1:2.8 lens, tripod, several spirit levels, Transcend[®] Compact Flash 8GB 133x memory card, custom made panoramic head, gray card (for white balancing), a graphics workstation class PC with the required software (as described in Chapter 3);
- Mirrored Ball: Canon[®] EOS 1Ds Mark II (17 megapixel), Canon[®] EF 70-300 mm 4: 5.6 zoom lens, 60 mm G100 mirrored ball, two tripods, several spirit levels, Transcend[®] Compact Flash 8GB 133x memory card, gray card (for white balancing), a graphics workstation class PC with the required software (as described in Chapter 3).

The costs for the fisheye and mirrored ball techniques refer to averaged prices on the European market available by the time of this writing.

In case of the SpheroCam[®] HDR the overall acquisition time and image quality can become a problem in low light environments if the camera must cover its full DR. The line CCD (Charge-Coupled Device) sensor manufactured by Sony and used in the SpheroCam[®] HDR (Sony ILX518K, discontinued product) has a pixel size of only 8 x 8µm (compare to ILX533K, which is a high-sensitivity version of Sony ILX518K [60]) which results in relatively noisy images at low illumination levels, especially when using long exposure times, due to heating of the sensor and the resulting thermal noise [59]. The acquisition time problem may arise if a high quality environmental map is required (e.g. to be used as a background image). As shown in Figure 5.3, for the otherwise identical camera parameters capturing the environment at the maximum resolution takes up twice as much time as at the “High” resolution setting of 6981 x 2624 pixels.

The fisheye technique (with the particular selection of hardware for this thesis) on the other hand does require certain post-processing steps but the acquisition time is essentially resolution independent (unless a slow memory card is used). Furthermore due to usage of a CMOS (Complementary metal-oxide-semiconductor) sensor and additional image optimization techniques [61] the Canon[®] EOS-1Ds Mark II produces low noise images even in low light environments.

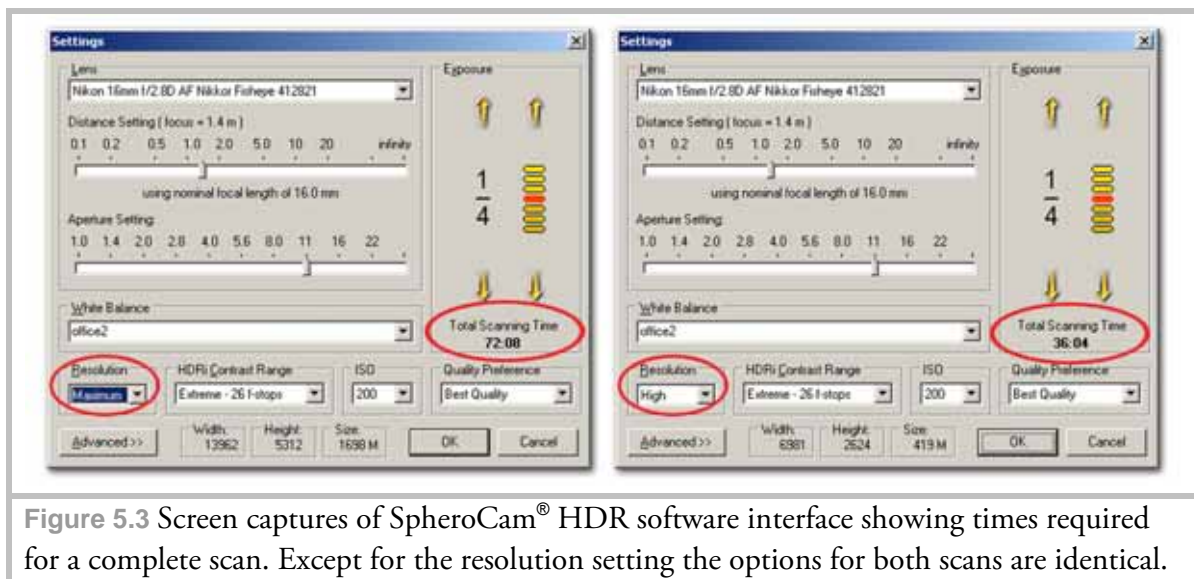


Figure 5.3 Screen captures of SpheroCam[®] HDR software interface showing times required for a complete scan. Except for the resolution setting the options for both scans are identical.

The mirrored ball technique offers the most inexpensive solution. Depending on the production demands a non-DSLR digital camera can be used, as many cameras nowadays offer RAW capabilities. In Europe prices for mirrored balls range from 20€ (yard balls) to 500€ (chrome coated precision bearing balls). One particular drawback of this technique is the fact that the photographer has to change his and camera's position during shooting while retaining the distance to the ball. On the other hand depending on the chosen ball a good quality HDR image can be obtained which is well suited for both lighting and reflection purposes.

	SpheroCam® HDR	Fisheye	Mirrored ball
Pure capturing time (not including the setup time, white balancing etc.)	about 8-13 minutes in an outdoor environment; up to 70 minutes in an indoor low light environment	about 12 minutes (depending on the photographic skills)	about 6 minutes (depending on the photographic skills)
Maximum resolution (pixels)	13962 x 5312	13346 x 6673	6000 x 3000
Post processing time (incl. file format conversions, HDR assembly, panoramic assembly)	none	up to two hours	up to an hour (or more depending on the mirrored ball and the operator's skills)
Skills required	<ul style="list-style-type: none"> - operating the SpheroCam® HDR and its software; - operating an image editing software (for retouching) 	<ul style="list-style-type: none"> - operating the selected digital camera; - processing the captured data; - operating a panoramic assembly software - operating an image editing software (for retouching) 	<ul style="list-style-type: none"> - operating the selected digital camera; - processing the captured data; - operating an image editing software (for retouching and combining of multiple views of the light probe)
Cost factor	about 55.000 €	about 10.000 €	about 6000-7000 €

Table 5.1 Comparison of techniques

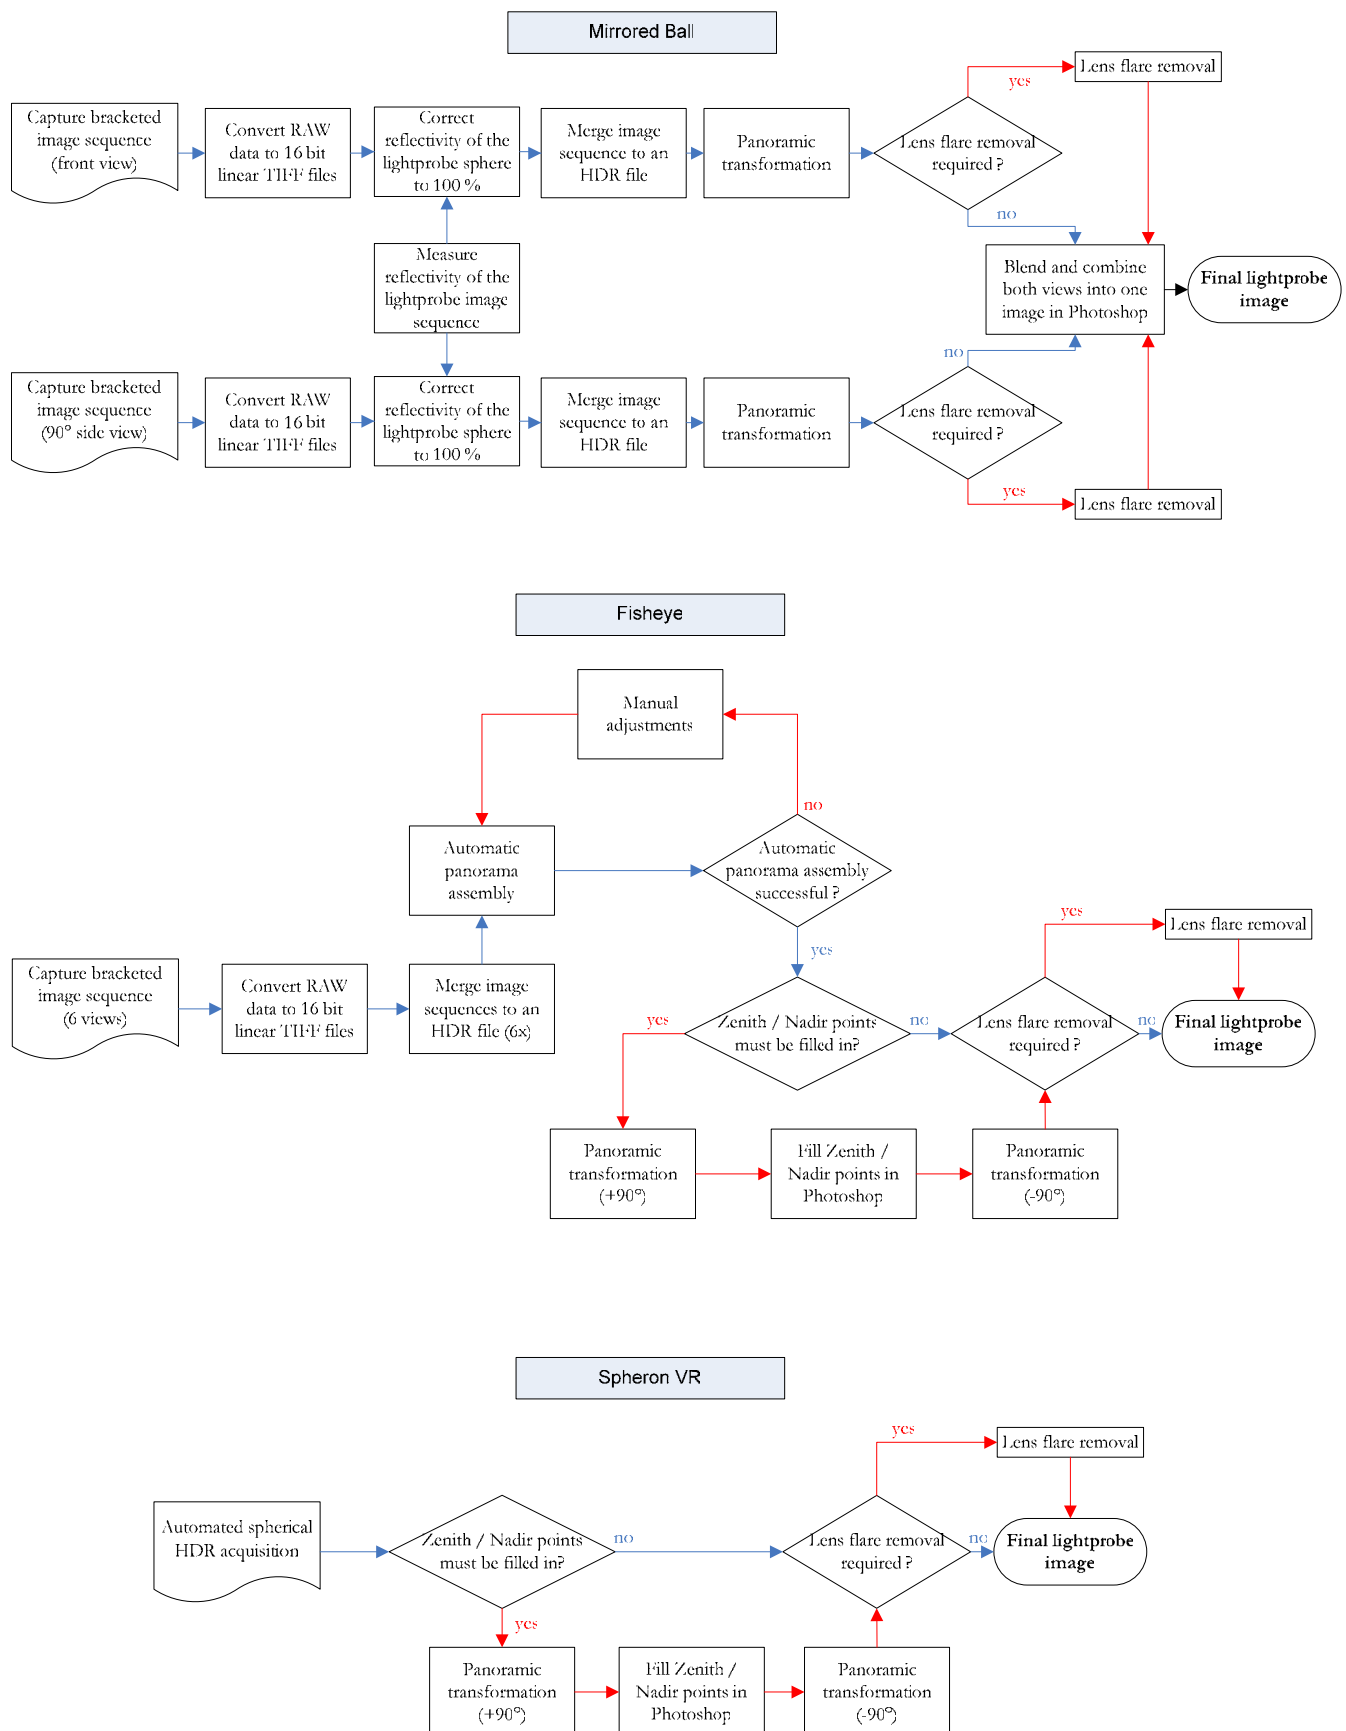


Figure 5.4 Comparison of techniques (continued)

**Figure 5.5 Results**

5.3 Dealing with light source intensities

Our review would not be complete without stressing the importance of light source intensities captured by a digital camera, especially that of the sun. Since the sun is the brightest natural light source and can be 50.000 times brighter than the sky [33], its intensity can not be captured even at the shortest exposure time and the smallest aperture setting of a digital camera, because even at these settings pixels of the sensor get saturated in the area covered by the sun. In IBL detailed shadows can be produced only by non-saturated light probe images. The SpheroCam® HDR is the only device capable of capturing the sun's intensity without any additional technical aids.

Several techniques [33] have been introduced to extract the color information of the saturated sun region. One technique makes use of additional ND (neutral density) filters in conjunction with certain exposure time and aperture combinations to allow proper non-sensor-saturating capture of high intensity light sources. Another solution derives the sun's color information indirectly from a matt white (or gray) ball photographed at the same location as the light probe. Since the matt ball reflects light diffusely the color information of the sun can be easily derived from its surface. An additional mathematical operation (as described in section 2.1) is applied to calculate the directional vector of the sun. This information can then be applied to a simulated artificial light source in a virtual scene or in an even more appropriate way to a physical sun simulation system which is available in many professional 3D applications.

Figure 5.6 shows a side-by-side comparison of shadows generated by an IBL application through “saturated” and “non-saturated” light probe images respectively, where the effect of soft shadows due to insufficiently covered DR becomes apparent.





Figure 5.6 Shadows produced by saturated (left) and non-saturated (right) light probe images

5.4 A typical HDRI workflow for virtual scene lighting



The assembled HDR environments can now be used to light an artificial scene. This technique is widely used in the CGI industry either in conjunction with additional virtual light sources, or to match lighting conditions in a virtual scene to a background photograph.

One commonly used HDR lighting technique involves projecting the equirectangular representation of a spherical HDR environment onto the surface of a virtual sphere. The lighting information is then derived from individual points on the inner side of the sphere which get sampled by the rendering software. The samples transfer color values of the environment (and thus the lighting color information) to the virtual objects of the scene (for instance blue color variations of the sky transfer to virtual objects where they become color tint information on the object's surface).

The sampling density determines the lighting quality and the rendering time. While high resolution environmental images may be required to produce clear reflections on mirror-type surfaces, sampling them for lighting purposes may not be beneficial, as with increasing number of samples this technique begins to produce grainy artifacts in the rendered image. In such situations it is best to create a low resolution copy of the environmental image (and even blur it in an image editing application) which is then used for lighting purposes only, while retaining the high quality version for reflection calculations. The sampling difference and the corresponding rendering results are shown in Figure 5.8. A software solution called "Smart IBL" [62] is capable of automatic generation of multiple versions of the same environmental image at different resolutions. Additional information on HDR light rendering algorithms can be found in [63].



Figure 5.8 Difference between densely (left) and coarsely (right) sampled image

5.5 Final results: Audi R8 CGI renditions

Figure 5.9 shows the final composited (as described in Chapter 4) renditions of the car model with different HDR environmental images applied for lighting and reflection. The original rendered image resolution equals the background image resolution of 17 megapixels. The rendering was accomplished using the Spheron® VR rendering module for Autodesk Maya.

As the interior lighting conditions shown here were mixed with the outdoor daylight, the overall lighting conditions shifted during the capture. Therefore the color matching process of the resulting images proved to be extremely difficult and had to be skipped. Nevertheless, as the main goal of the test was to compare the reflections and shadows produced by the different environmental images the color differences did not contribute negatively to these image features.

It was determined that both fisheye and mirrored ball techniques (having the SpheroCam® HDR image as the reference) produced crisp shadows without artifacts. The reflections produced by the mirrored ball image are not as clear as those of the other two techniques, though.



Figure 5.9 (from left to right) Rendering results of the car model in Autodesk Maya: mirrored ball, fisheye, SpheroCam® HDR (reproduced with permission from 809cgi GbR)

5.6 Problems of the fisheye technique and a proposal for their solution

Comparing the fisheye results with those of the SpheroCam® HDR we can say that the quality, resolution and cost factor of this technique are superior. However, there are two major drawbacks, namely the image acquisition and processing times. These make the technique unsuitable for on-location shooting, where the HDR results are required shortly after the acquisition step to present combined CGI preview results to a customer. To resolve the time consumption problem, apart from using an automated panoramic head (proposed earlier in the conclusion of Chapter 3) a software application could be developed to automate the panorama generation step. An idea of a suitable GUI and a processing algorithm is shown in Figure 5.10. The user would be able to control several aspects of the process, such as processing of zenith / nadir points, or additional options of the RAW data conversion. The software would then automatically sort image sequences (for instance by creation time or image number), process the converted RAW data to individual HDR images of the respective view, and finally generate an equirectangular HDR panorama. The benefit of such an automated solution would be the processing time, as the user would have the option to specify the resolution for the preview image of the final HDR panorama. For instance, if for a 13346 x 6673 pixel full resolution image a 2000 x 1000 pixel preview were specified (which is sufficient for the most preview purposes) the processing (depending on the computing power available) would take far less time. An instant preview feature (i.e. visual feedback of the current processing step) might also prove useful, e.g. to display a tone-mapped version of the HDR panorama inside the proposed tool.

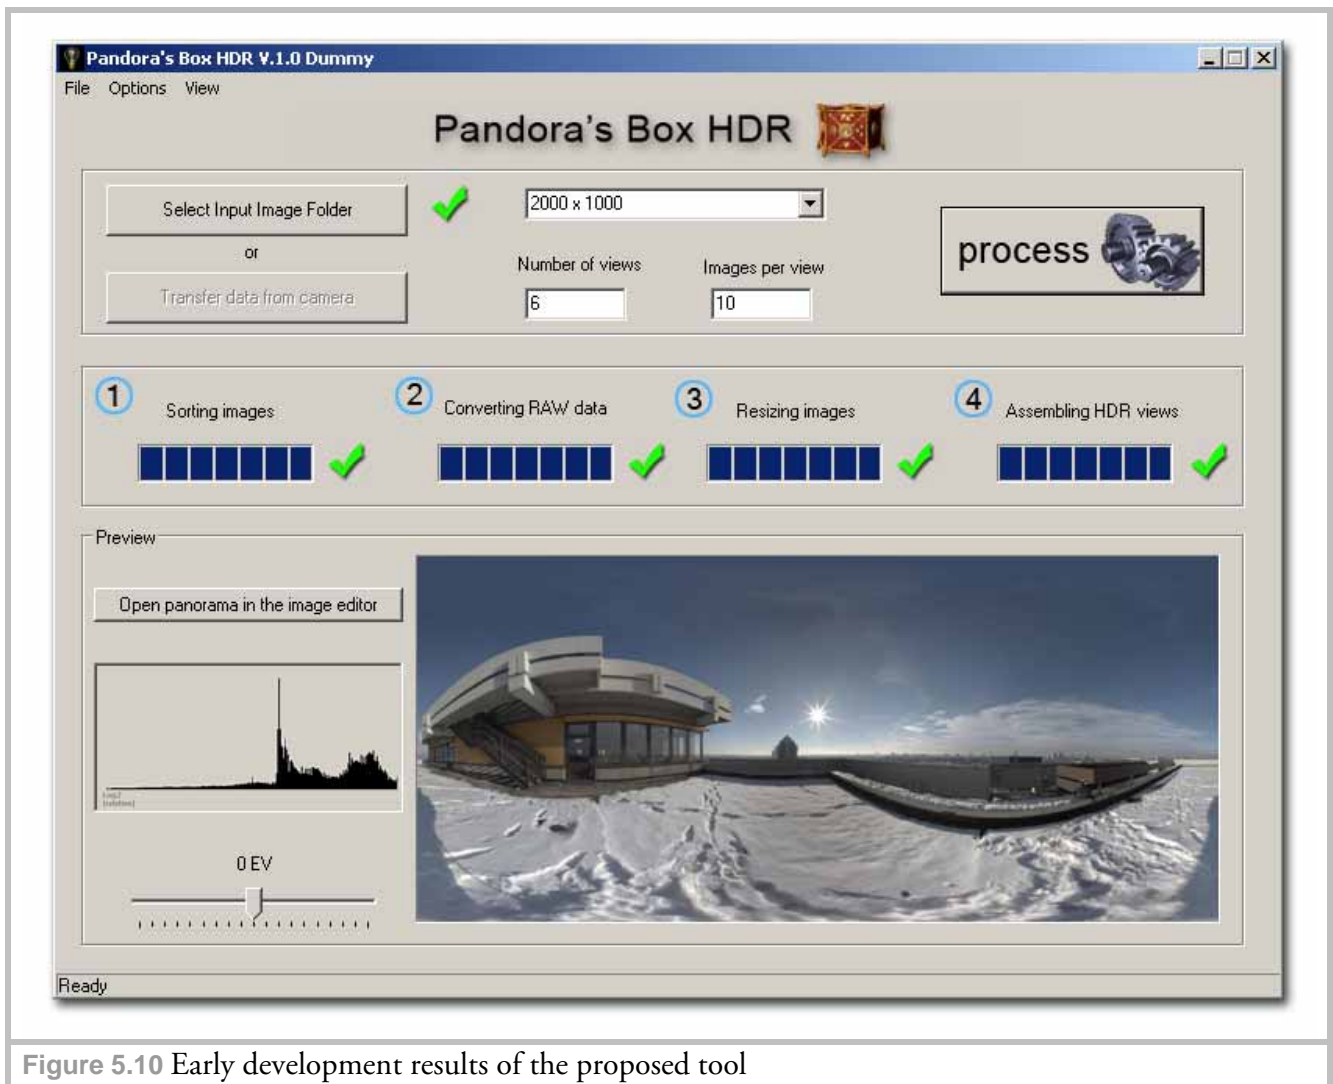


Figure 5.10 Early development results of the proposed tool

An additional advantage of such automatic processing would be that, once satisfied with the result, the user could process the original full resolution images to produce the high quality HDR environmental image.

The user interface proposed could link several freeware programs together to produce the final image, but a more robust solution would be desirable. The main processing module of the software would be the panoramic assembly algorithms and for this purpose the Panorama Tools software library (<http://panotools.sourceforge.net/>) is well suited. Most of the low cost commercial panoramic software packages (like PTGui® Pro and Autopano® Pro) are based on this library. Furthermore a variety of freeware tools for individual steps (like HDR assembly with the “PFScalibration” tool [64]) are available.

Conclusion and final thoughts

We have shown that both fisheye and mirrored ball techniques are legitimate alternatives to high-end professional solutions like SpheroCam[®] HDR. In terms of productivity these techniques take up more time but are inexpensive compared to their high-end competitor.

High quality reflection images are required in professional high resolution still photography and for special effects during movie productions. On the other hand, there are many applications (e.g. broadcast TV or video games) where the mirrored ball image quality would be sufficient. If lighting conditions are predictable (e.g. using controllable light sources) the fisheye technique may produce even better results than the high-end competitor SpheroCam[®] HDR.

Due to the fact that digital cameras nowadays can cover only a limited DR a cost-effective way of capturing high intensity light sources is required.

Of particular interest may also be the gamut comparison between the SpheroCam[®] HDR and the Mark II camera.

Many companies today offer light probe acquisition services delivering high quality HDR images at a relatively low price. One of the best known is Dosch Design [65] who offer each of their HDRI bundles at a price of 99 € at the time of this writing. There are many situations in which an HDR environmental image serves only as indirect illumination light source to provide an object with a more pleasing look. Therefore a matching photograph of the background is not required and companies tend to buy prefabricated images which can be used over and over again. For studio use, lighting techniques are available to generate HDR environmental images through a software solution (like Autodesk[®] Maya[®]). The different application areas demand light probe images of different quality which gives the visual effects production companies the ability to choose a suitable solution.

References:

- [1] Debevec, Paul. *Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-based Graphics with Global Illumination and High Dynamic Range Photography*. SIGGRAPH 98, July 1998.
- [2] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 293, Wiley Publishing: Indianapolis, Indiana, 2008.
- [3] Palmer, James M. *Radiometry and photometry FAQ*: Optical Sciences Center University of Arizona,
<http://www.optics.arizona.edu/Palmer/rpfaq/rpfaq.pdf>.
- [4] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 19, Morgan Kaufmann: San Francisco, 2005.
- [5] <http://www.bipm.org/en/si/>.
- [6] http://web.cs.wpi.edu/~emmanuel/courses/cs563/write_ups/chuckm/chuckm_BRDFs_overview.html.
- [7] Commerce, U.S. Department of. *Guide for the Use of the International System of Units (SI)*. NIST Special Publication (SP) 811, National Institute of Standards and Technology: Gaithersburg, 1995 Edition, MD 20899.
- [8] http://www.schorsch.com/kbase/glossary/solid_angle.html.
- [9] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 24, Morgan Kaufmann: San Francisco, 2005.
- [10] http://upload.wikimedia.org/wikipedia/commons/7/72/CIE_1931_Luminosity.png.
- [11] Glenn M. Rand, David R. Litschel, Robert G. Davis. *Digital Photographic Capture*, p. 181, Focal Press, 2005.
- [12] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 76, Morgan Kaufmann: San Francisco, 2005.

- [13] Bloch, Christian. *The HDRI Handbook: High Dynamic Range Imaging for Photographers and CG Artists*, p. 15, Rocky Nook, 2007.
- [14] Kerr, Douglas A. *APEX—The Additive System of Photographic Exposure* p. 6, <http://doug.kerr.home.att.net/pumpkin/APEX.pdf>.
- [15] Glenn M. Rand, David R. Litschel, Robert G. Davis. *Digital Photographic Capture*, p. 54, Focal Press, 2005.
- [16] Bloch, Christian. *The HDRI Handbook: High Dynamic Range Imaging for Photographers and CG Artists*, p. 18, Rocky Nook, 2007.
- [17] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 136-147, Morgan Kaufmann: San Francisco, 2005.
- [18] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 136, Morgan Kaufmann: San Francisco, 2005.
- [19] Fraser, Bruce. *Understanding Digital Raw Capture*, p. 1, http://www.adobe.com/digitalimag/pdfs/understanding_digitalrawcapture.pdf.
- [20] Fraser, Bruce. *Understanding Digital Raw Capture*, p. 2, http://www.adobe.com/digitalimag/pdfs/understanding_digitalrawcapture.pdf.
- [21] Géraldine Joffre, William Puech, Frédéric Comby, Jacques Joffre. *High Dynamic Range Images from Digital Cameras Raw Data*, www.hdrsoft.com/resources/hdri_from_raws_poster_abstract.pdf, 2005.
- [22] Géraldine Joffre, William Puech, Frédéric Comby, Jacques Joffre. *High Dynamic Range Images from Digital Cameras Raw Data (Poster)*, http://www.hdrsoft.com/resources/HDRIfromRAWs_posterA4.pdf, 2005.
- [23] http://www.drycreekphoto.com/tools/printer_gamuts/gamutmodel.html.
- [24] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 89-90, Morgan Kaufmann: San Francisco, 2005.
- [25] Ward, Greg. *High Dynamic Range Image Encodings*, <http://www.anywhere.com/gward/hdrenc/Encodings.pdf>.

- [26] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 381-382, Morgan Kaufmann: San Francisco, 2005.
- [27] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 402, Morgan Kaufmann: San Francisco, 2005.
- [28] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 405, Morgan Kaufmann: San Francisco, 2005.
- [29] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 383, Morgan Kaufmann: San Francisco, 2005.
- [30] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 48-49, Morgan Kaufmann: San Francisco, 2005.
- [31] Johnson, Dave. *How to Do Everything: Digital Camera* p. 111-112, McGraw-Hill Osborne Media, 2008.
- [32] http://www.hdrsoft.com/support/faq_photomatix.html#exposure.
- [33] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 392-401, Morgan Kaufmann: San Francisco, 2005.
- [34] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 152-160, Morgan Kaufmann: San Francisco, 2005.
- [35] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 383-385, Morgan Kaufmann: San Francisco, 2005.
- [36] *Wälzlager - Kugeln für Wälzlager und allgemeinen Industriebedarf*, p. 4, DEUTSCHE NORM, DIN 5401, August 2002.
- [37] http://wiki.panotools.org/Fisheye_Projection.
- [38] Greene. *Creating Raster Omnimax Images from Multiple Perspective Views Using the Elliptically Weighted Average Filter*, p. 21-27: IEEE Computer Graphics and Applications, June 1986.

- [39] Max, Nelson. *Computer Graphics Distortion for IMAX and OMNIMAX Projection*, p. 137-159: Proc. Nicograph 83, Dec. 1983.
- [40] http://www.bobatkins.com/photography/technical/field_of_view.html.
- [41] http://www.canon.de/For_Home/Product_Finder/Cameras/EF_Lenses/Fixed_Focal_Length/EF_15mm_f28_Fisheye/index.asp.
- [42] Bloch, Christian. *The HDRI Handbook: High Dynamic Range Imaging for Photographers and CG Artists*, p. 264-265, Rocky Nook, 2007.
- [43] Bloch, Christian. *The HDRI Handbook: High Dynamic Range Imaging for Photographers and CG Artists*, p. 104, Rocky Nook, 2007.
- [44] Brinkmann, Ron. *The Art and Science of Digital Compositing (The Morgan Kaufmann Series in Computer Graphics)*, p. 2-3, Morgan Kaufmann, 1999.
- [45] Wright, Steve. *Compositing Visual Effects: Essentials for the Aspiring Artist*, p. 45, Focal Press, 2007.
- [46] *The University of Edinburgh Informatics Homepages Server*, <http://homepages.inf.ed.ac.uk/rbf/HIPR2/pixadd.htm>.
- [47] Wright, Steve. *Digital Compositing for Film and Video (Focal Press Visual Effects and Animation Series)*, p. 97, Oxford: Elsevier LTD, 2001.
- [48] Birn, Jeremy. *Digital Lighting & Rendering, Second Edition*, Chapter eleven: Rendering Passes and Compositing, New Riders, 2006
- [49] <http://forensicphotoshop.blogspot.com/2007/09/blending-mode-math-part-2.html>.
- [50] Lanier, Lee. *Maya : professional tips and techniques* p. 157, Indianapolis, Indiana: Wiley Publishing, 2007.
- [51] Phong, Bui Tuong, "*Illumination for Computer Generated Pictures*" Communications of the ACM, 18(6): p. 311-317, June 1975.
- [52] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 622, Wiley Publishing: Indianapolis, Indiana, 2008.
- [53] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 625, Wiley Publishing: Indianapolis, Indiana, 2008.

- [54] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 626, Wiley Publishing: Indianapolis, Indiana, 2008.
- [55] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 70, Wiley Publishing: Indianapolis, Indiana, 2008.
- [56] Livny, Boaz. *Mental Ray for Maya, 3ds Max, and XSI*, p. 71, Wiley Publishing: Indianapolis, Indiana, 2008.
- [57] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 11-12, Morgan Kaufmann: San Francisco, 2005.
- [58] <http://www.spheron.com/en/intruvision/solutions/spherocam-hdr/technical-details.html>.
- [59] <http://www.panoscan.com/Advantage/index.html>.
- [60] http://www.sony.net/Products/SC-HP/cx_news/vol09/pdf/ilx533.pdf.
- [61] http://www.canon.de/For_Home/Product_Finder/Cameras/Digital_SLR/EOS_1Ds_markII.
- [62] <http://www.hdrlabs.com/sibl/index.html>.
- [63] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, p. 409-414, Morgan Kaufmann: San Francisco, 2005.
- [64] <http://www.mpi-inf.mpg.de/resources/hdr/calibration/pfs.html>.
- [65] <http://www.doschdesign.com/products/hdri/>.

APPENDIX A

Usage of DCRAW GUI

DCRAW is a freeware RAW format converter for Microsoft® Windows® created by Dave Coffin (<http://www.cybercom.net/~dcoffin/>). Apart from the fact that it supports most of the RAW file formats available today (for both amateur and professional cameras), it also has some features of particular interest to HDR imaging that are not available (or hardly accessible) in other RAW converters. As DCRAW is a command line tool, converting a group of images requires a lot of manual command input and is therefore best accomplished via batch processing. As normally all images for HDR generation are processed with the same settings a graphical user interface may be of great help for repetitive operations (e.g. processing individual views of fisheye image sets). Such an interface, called “DCRAW GUI” was developed during this thesis.

The user interface is divided into two areas: the “Main Window” where images are loaded and saved and the “Options” area, where individual RAW processing features are selected. Please note that this interface was developed only for the purpose of HDR image assembly and currently does not support many features available in DCRAW. There are also some usability limitations which are discussed in the next section. Future development is planned, as well as publishing under the GNU General Public License.

The “Main Window” area

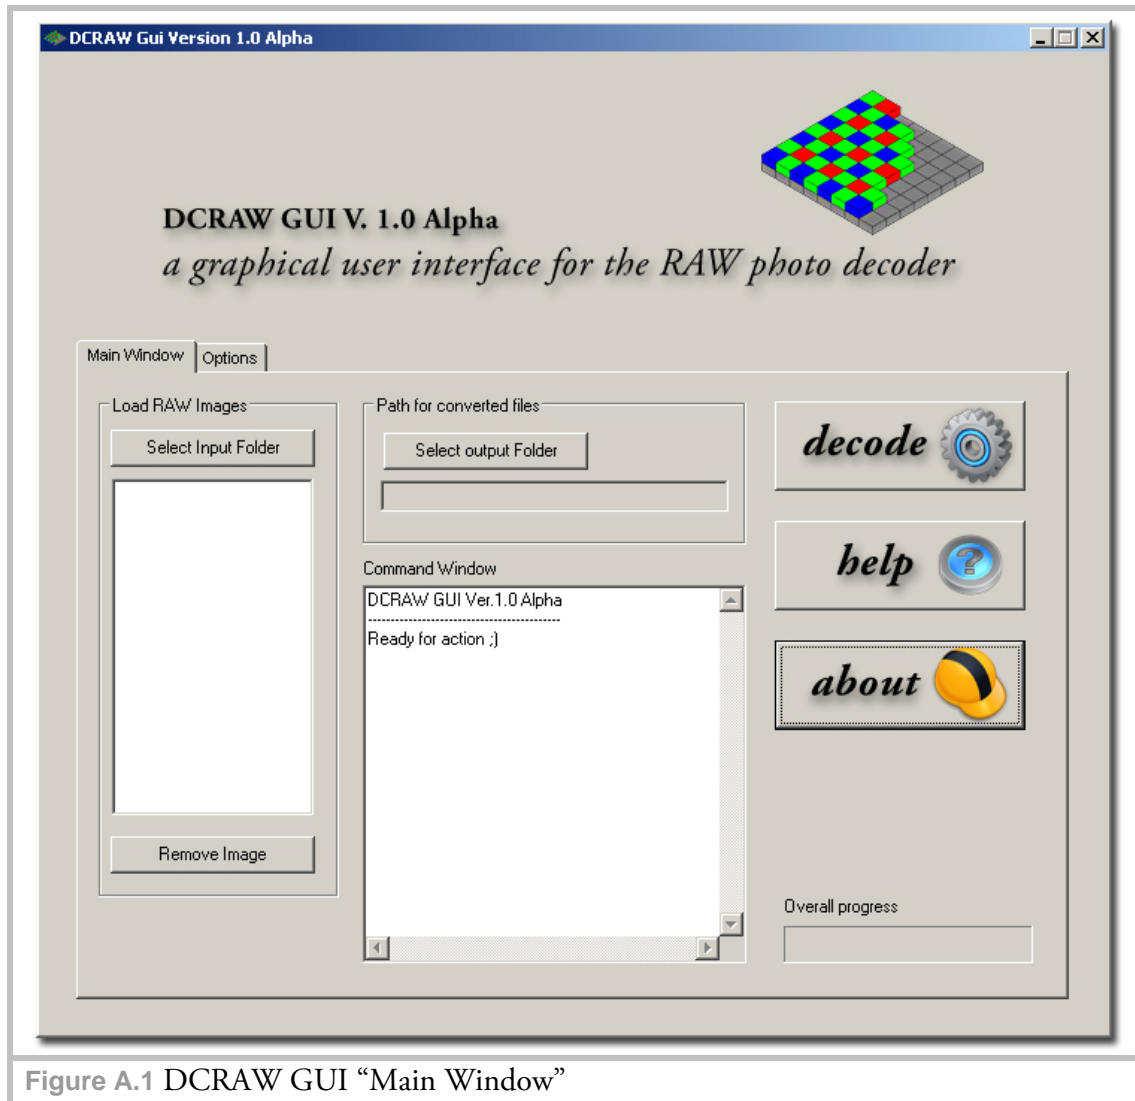
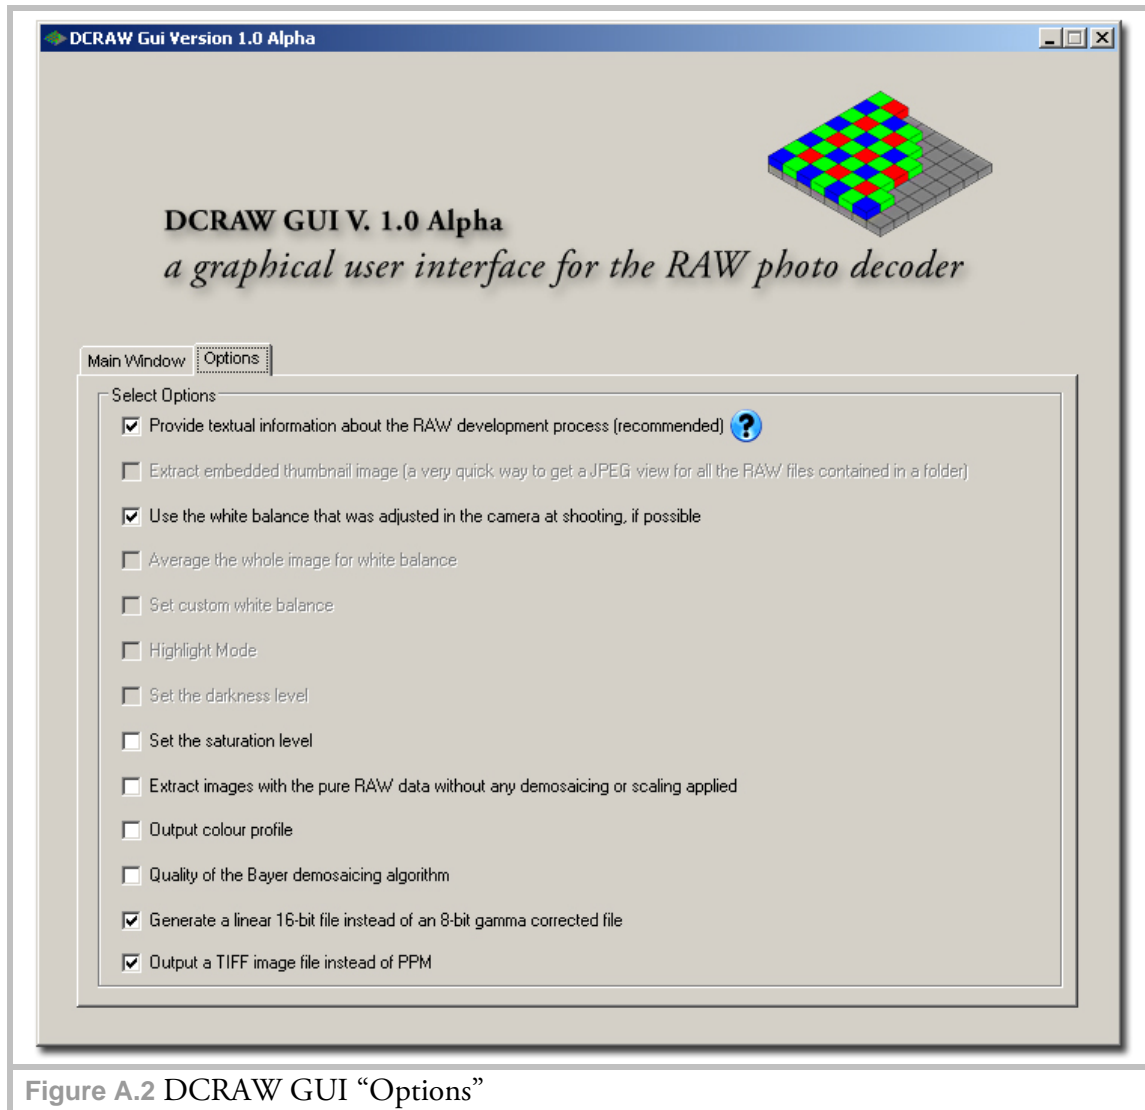


Figure A.1 DCRAW GUI “Main Window”

Controls in this area are self-explanatory. There are following limitations to be aware of in the current version:

- if the Input Folder contains subfolders these are added to the processing queue as well. In this case the subfolder entries should be removed from the list with the “Remove Image” button;
- images, files or folders can be removed from the queue only one at a time – multiple selection is currently not supported;
- both “Select Input Folder” and “Select Output Folder” do not remember the last specified file location.

The “Options” area



The description of individual options has been taken from (http://www.guillermoluijk.com/tutorial/DCRAW/index_en.htm), where additional information on all features can be found. The ? icon provides additional information for a particular option (currently under development). For mirrored ball measurements the most important setting is the generation of 16-bit linear data which yields radiometrically correct images.

For HDRI purposes the following options retain most of the valuable sensor data:

- linear 16-bit file generation;
- TIFF output for processing in Photomatix® Pro;

- RAW color profile (if no output color profile is specified the default profile used is sRGB).

Saturation Level

Under certain circumstances DCRAW may not have proper saturation information for a particular camera. In such cases it provides a way to extract pure RAW data from the input imagery which allows us to determine the brightness level at which the camera's sensor gets saturated (option: "Extract images with the pure RAW data without any demosaicing or scaling applied").

After the conversion step the image should be loaded in any application capable of displaying

RGB histogram plots, in order to provide the necessary information on the range of levels spanned by the dynamic range of the image. A good freeware alternative is "Histogrammar" (<http://www.guillermoluijk.com/software/histogrammar/index.htm>) which displays the required information directly as a numeric value (Figure A.3). The values displayed are dependent on the bit depth of the sensor (i.e. the higher the bit depth of the sensor the greater the values may be). After the value has been determined it can be entered into the "Set the saturation level" field of DCRAW GUI (not available in the current version, shall be implemented later). For the measurements and panoramic image acquisition in this thesis the camera's saturation level was tested (Canon® EOS 1Ds Mark II). On a side note, the built-in values of DCRAW proved to be correct for this particular camera (DCRAW developed all the images at 3712).

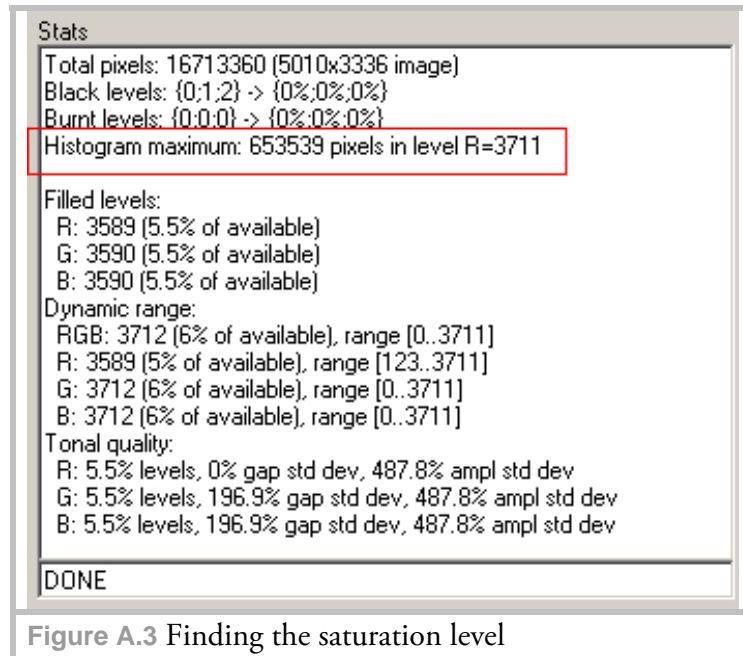


Figure A.3 Finding the saturation level

Appendix B

Using the MATLAB[®] tools

In the following section we shall discuss two MATLAB[®] programs created specifically to aid in the mirrored ball HDR image acquisition process. Please note that due to MATLAB[®] licensing policies only the source code for the modules can be provided.

B.1 Reflectance calculation tool

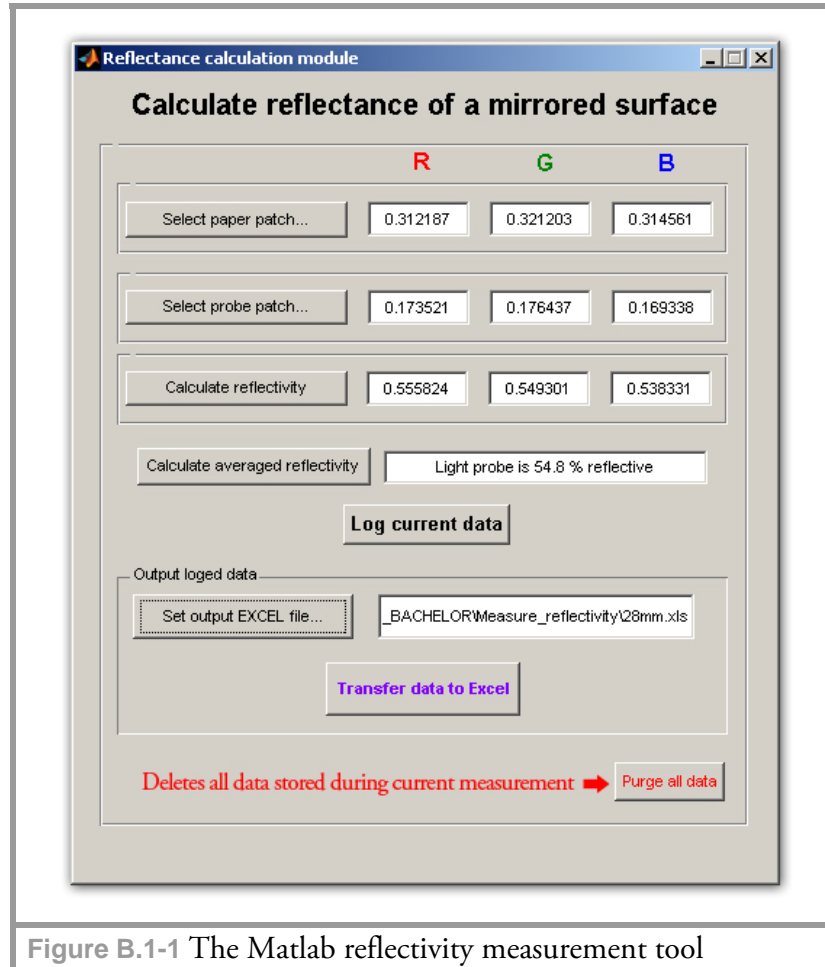


Figure B.1-1 The Matlab reflectivity measurement tool

This tool allows to measure relative reflectivity values of a mirrored sphere for its individual color channels. The data acquired can be temporarily stored and then transferred to a Microsoft® Excel® spreadsheet specified in the “Output logged data” area.

The images of the patches (described in Chapter 2) should be placed in the root directory of the compiled version of the reflectance calculation module.

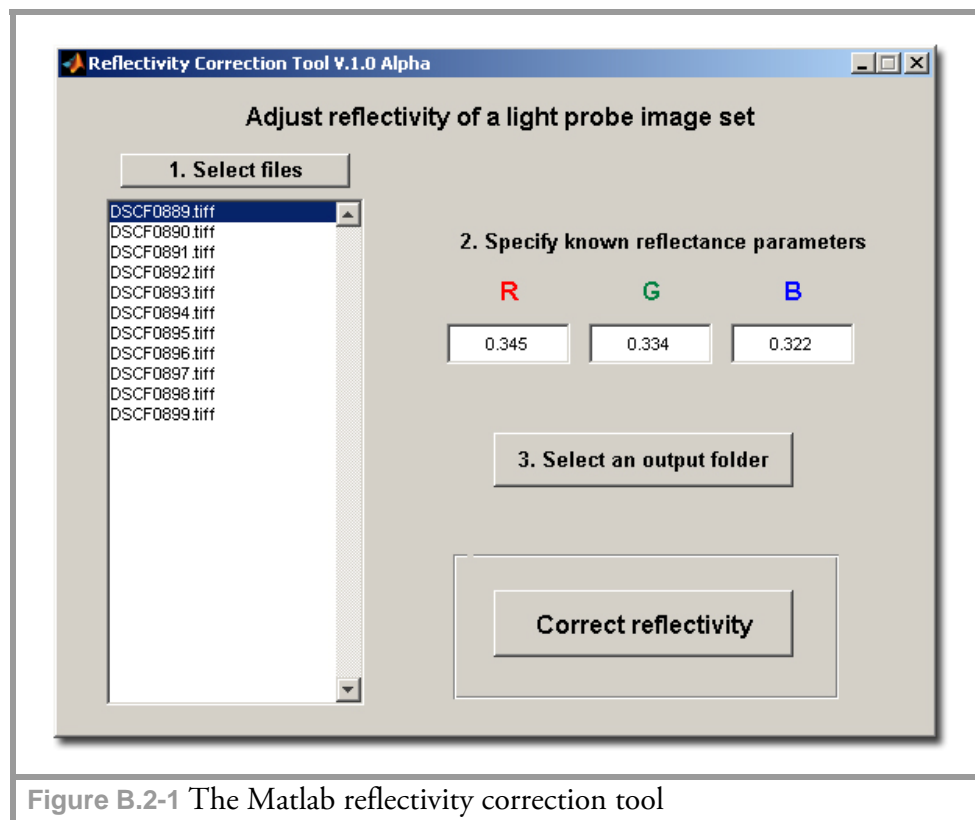
After both patches have been selected we can calculate the reflectivity for each individual color channel (“Calculate reflectivity” button) or the overall reflectivity of the light probe (mean of all color channels). Then data can be logged and additional patches can be measured, each time logging previous values. When all values of a measurement set have been acquired we can set an output Excel® file where the values will be stored. If necessary, all temporarily stored data can be deleted with the “Purge all data” option.

The spreadsheet template has been prepared specifically to process the transferred data. When the tool stores measured data in a temporary file the data of the previous measurement is preserved on the next measurement. The data transferred to the Excel® spreadsheet always begins at cell A3. If a single reflectivity measurement is to be made, the way the measurement tool saves the data is not an issue. If multiple measurements are to be made, the user has two options to choose from:

- 1) All the measurements are made at once without closing the application. New data is stored after the previous one in a single Excel® spreadsheet, e.g. if the first measurement spans from A3 to A11, then the second measurement would start at cell A12.
- 2) If measurements are to be made separately and the tool needs to be closed and re-opened in the process, several copies of the template spreadsheet should be made (e.g. ball1.xls, ball2.xls etc.), as each new measurement would then require its own spreadsheet.

If the same spreadsheet were used after restarting the tool, the data transfer would start over again at cell A3, overwriting existing data of the previous measurements.

B.2 The reflectivity correction module



Having calculated the relative reflectivity of a mirrored ball we can now apply these values to correct individual mirrored ball images.

Usage of the reflectivity correction module is self-explanatory except for the following limitations in the current version:

- the *EXIF data* (metadata inside the image containing exposure time, focal length etc.) is not preserved throughout the correction process. As HDR image assembly tools (such as Photomatix® Pro) frequently rely on the exposure time and the aperture data, this data has to be assigned manually during the assembly (as shown in Figure B.3 for Photomatix® Pro). Therefore it is recommended to properly rename the files prior to HDR image creation, giving each image a name prefix “EV” followed by a positive / negative numeric value;
- last used Input / Output folders are not stored – the default directory is the root folder of the compiled version of the tool;
- more than one file has to be selected for processing – single file correction is not supported yet.

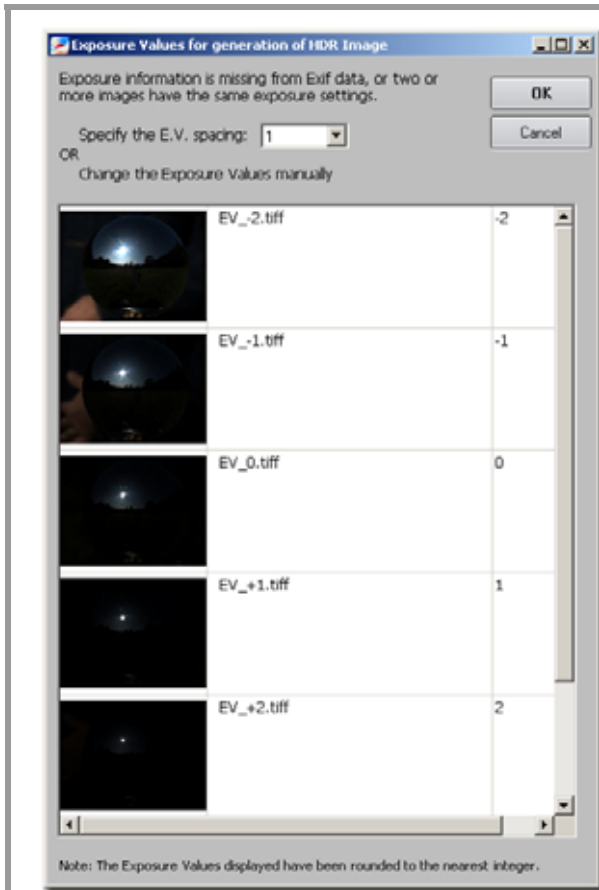


Figure B.2-2 The Matlab® reflectivity correction tool



Figure B.2-3 Image comparison (corrected vs. unprocessed)

The tool supports the images in TIFF and JPEG file formats. The corrected images have the prefix “calibrated_.*”.

The processed sequence can then be assembled to an HDR image which will be radiometrically consistent with a *backplate* (photograph of a particular area of interest taken on the same location).

B.3 Reflectivity measurements of randomly selected mirrored spheres

Measurement No. (N)	Paper patch			Probe patch			Calculated Reflectivity			averaged reflectivity $\bar{\rho}_{\text{sample (N)}}$ decimal value	overall averaged reflectivity $\bar{\rho}_{\text{sphere}}$	standard deviation σ_{sphere}	Sphere type
	R	G	B	R	G	B	R	G	B				
1	0.606	0.615	0.658	0.340	0.340	0.352	0.560	0.552	0.535	0.549	55.0 %	± 0.46%	28 mm bearing ball
2	0.596	0.604	0.644	0.335	0.335	0.346	0.562	0.554	0.538	0.551			
3	0.621	0.629	0.674	0.351	0.352	0.367	0.566	0.559	0.544	0.556			
4	0.625	0.634	0.681	0.352	0.353	0.369	0.564	0.557	0.541	0.554			
5	0.612	0.620	0.664	0.340	0.339	0.351	0.555	0.547	0.528	0.543			
6	0.610	0.618	0.662	0.342	0.342	0.355	0.561	0.553	0.536	0.550			
7	0.607	0.616	0.659	0.344	0.345	0.359	0.567	0.560	0.545	0.557			
8	0.622	0.632	0.678	0.347	0.347	0.361	0.558	0.550	0.533	0.547			
9	0.600	0.608	0.650	0.335	0.334	0.345	0.557	0.549	0.530	0.545			
10	0.611	0.619	0.663	0.342	0.342	0.354	0.560	0.551	0.534	0.548			
11	0.481	0.482	0.495	0.264	0.260	0.256	0.550	0.539	0.518	0.536	53.7 %	± 0.85%	35 mm bearing ball
12	0.500	0.500	0.514	0.277	0.273	0.270	0.555	0.545	0.525	0.542			
13	0.505	0.507	0.522	0.283	0.280	0.279	0.561	0.552	0.534	0.549			
14	0.505	0.507	0.522	0.279	0.275	0.272	0.553	0.542	0.521	0.539			
15	0.503	0.505	0.521	0.276	0.272	0.269	0.548	0.538	0.517	0.535			
16	0.516	0.518	0.537	0.285	0.281	0.279	0.552	0.541	0.520	0.538			
17	0.476	0.478	0.491	0.267	0.264	0.262	0.560	0.551	0.533	0.548			
18	0.498	0.501	0.515	0.267	0.262	0.257	0.535	0.522	0.500	0.519			
19	0.511	0.513	0.530	0.280	0.275	0.273	0.547	0.535	0.514	0.532			
20	0.475	0.478	0.490	0.262	0.258	0.254	0.551	0.540	0.518	0.536			
21	0.836	0.861	0.927	0.451	0.459	0.480	0.539	0.534	0.518	0.530	52.6 %	± 0.68%	45 mm bearing ball
22	0.835	0.860	0.926	0.449	0.457	0.477	0.537	0.531	0.515	0.528			
23	0.906	0.928	0.979	0.484	0.493	0.514	0.535	0.532	0.525	0.530			
24	0.862	0.886	0.949	0.458	0.466	0.485	0.531	0.526	0.511	0.523			
25	0.827	0.852	0.919	0.447	0.456	0.477	0.540	0.536	0.520	0.532			
26	0.884	0.907	0.967	0.476	0.486	0.507	0.539	0.536	0.525	0.533			
27	0.822	0.847	0.914	0.439	0.447	0.466	0.534	0.528	0.510	0.524			
28	0.802	0.826	0.894	0.433	0.442	0.462	0.540	0.535	0.517	0.531			
29	0.820	0.849	0.913	0.426	0.437	0.454	0.520	0.515	0.497	0.511			
30	0.801	0.830	0.894	0.426	0.437	0.454	0.532	0.527	0.508	0.522			
31	0.706	0.713	0.770	0.444	0.462	0.510	0.629	0.648	0.662	0.646	64.6 %	± 0.95%	45 mm Chinese baoding ball
32	0.699	0.708	0.761	0.437	0.455	0.499	0.625	0.643	0.656	0.641			
33	0.728	0.737	0.792	0.457	0.476	0.520	0.628	0.645	0.656	0.643			
34	0.658	0.667	0.717	0.425	0.444	0.488	0.645	0.665	0.681	0.664			
35	0.634	0.643	0.690	0.399	0.417	0.456	0.630	0.648	0.661	0.646			
36	0.676	0.685	0.736	0.419	0.437	0.478	0.620	0.637	0.650	0.636			
37	0.670	0.679	0.729	0.429	0.447	0.491	0.640	0.659	0.674	0.657			
38	0.722	0.731	0.785	0.456	0.475	0.520	0.632	0.650	0.662	0.648			
39	0.686	0.694	0.746	0.423	0.440	0.481	0.618	0.634	0.644	0.632			
40	0.707	0.716	0.770	0.448	0.468	0.514	0.634	0.653	0.667	0.652			
41	0.624	0.657	0.691	0.328	0.345	0.355	0.526	0.525	0.513	0.521	52.3 %	± 0.48%	55 mm bearing ball
42	0.644	0.678	0.714	0.339	0.356	0.366	0.527	0.525	0.513	0.522			
43	0.601	0.632	0.664	0.320	0.337	0.347	0.533	0.533	0.523	0.530			
44	0.602	0.634	0.665	0.318	0.333	0.342	0.527	0.526	0.514	0.522			
45	0.624	0.657	0.691	0.330	0.347	0.357	0.529	0.528	0.517	0.525			
46	0.624	0.658	0.692	0.330	0.347	0.358	0.529	0.528	0.517	0.525			
47	0.617	0.651	0.684	0.327	0.344	0.355	0.530	0.529	0.518	0.526			
48	0.607	0.639	0.671	0.322	0.338	0.348	0.531	0.530	0.519	0.527			
49	0.650	0.686	0.723	0.339	0.357	0.367	0.522	0.520	0.507	0.516			
50	0.580	0.610	0.639	0.302	0.316	0.322	0.520	0.518	0.504	0.514			

B.3 Reflectivity measurements of randomly selected mirrored spheres

Measurement No. (N)	Paper patch			Probe patch			Calculated Reflectivity			averaged reflectivity $\bar{\rho}_{\text{sample (N)}}$		overall averaged reflectivity $\bar{\rho}_{\text{sphere}}$	standard deviation σ_{sphere}	Sphere type
	R	G	B	R	G	B	R	G	B	decimal value	%			
51	0.679	0.687	0.739	0.413	0.427	0.462	0.608	0.621	0.626	0.618	61.8	62.1 %	± 1.46%	55 mm Chinese baoding ball
52	0.646	0.653	0.699	0.403	0.418	0.457	0.623	0.641	0.653	0.639	63.9			
53	0.695	0.703	0.755	0.417	0.432	0.470	0.600	0.614	0.622	0.612	61.2			
54	0.711	0.719	0.774	0.448	0.467	0.512	0.631	0.649	0.662	0.647	64.7			
55	0.705	0.713	0.766	0.413	0.428	0.468	0.585	0.600	0.611	0.599	59.9			
56	0.699	0.708	0.760	0.432	0.449	0.490	0.617	0.634	0.645	0.632	63.2			
57	0.683	0.691	0.742	0.406	0.422	0.462	0.595	0.611	0.622	0.609	60.9			
58	0.681	0.689	0.740	0.408	0.424	0.464	0.599	0.616	0.627	0.614	61.4			
59	0.700	0.708	0.761	0.423	0.438	0.476	0.604	0.618	0.626	0.616	61.6			
60	0.692	0.700	0.752	0.419	0.435	0.477	0.605	0.622	0.634	0.620	62.0			
61	0.607	0.617	0.663	0.334	0.338	0.353	0.550	0.548	0.533	0.544	54.4	54.2 %	± 0.18%	60 mm bearing ball
62	0.692	0.705	0.760	0.381	0.387	0.405	0.551	0.549	0.533	0.544	54.4			
63	0.628	0.640	0.690	0.345	0.350	0.366	0.550	0.547	0.531	0.542	54.2			
64	0.619	0.631	0.679	0.340	0.345	0.360	0.549	0.547	0.530	0.542	54.2			
65	0.625	0.637	0.685	0.344	0.349	0.365	0.551	0.549	0.533	0.544	54.4			
66	0.650	0.662	0.713	0.358	0.363	0.380	0.551	0.549	0.533	0.544	54.4			
67	0.644	0.656	0.706	0.352	0.357	0.372	0.547	0.544	0.527	0.539	53.9			
68	0.628	0.640	0.689	0.344	0.348	0.364	0.547	0.544	0.529	0.540	54.0			
69	0.650	0.662	0.713	0.357	0.362	0.378	0.549	0.546	0.530	0.542	54.2			
70	0.648	0.656	0.710	0.355	0.357	0.376	0.548	0.545	0.530	0.541	54.1			
71	0.511	0.519	0.548	0.294	0.297	0.305	0.575	0.571	0.556	0.567	56.7	56.4 %	± 0.73%	65 mm bearing ball
72	0.551	0.560	0.598	0.307	0.309	0.319	0.558	0.552	0.534	0.548	54.8			
73	0.525	0.532	0.566	0.302	0.305	0.317	0.576	0.574	0.559	0.570	57.0			
74	0.531	0.541	0.576	0.306	0.310	0.322	0.576	0.573	0.558	0.569	56.9			
75	0.509	0.515	0.544	0.294	0.296	0.305	0.577	0.574	0.560	0.570	57.0			
76	0.538	0.548	0.584	0.305	0.309	0.319	0.568	0.563	0.547	0.559	55.9			
77	0.519	0.528	0.560	0.299	0.303	0.313	0.576	0.573	0.559	0.570	57.0			
78	0.538	0.548	0.583	0.309	0.313	0.324	0.574	0.571	0.556	0.567	56.7			
79	0.525	0.535	0.568	0.297	0.301	0.311	0.566	0.563	0.549	0.559	55.9			
80	0.493	0.501	0.527	0.279	0.282	0.289	0.566	0.563	0.548	0.559	55.9			
81	0.532	0.544	0.577	0.283	0.288	0.296	0.532	0.529	0.512	0.524	52.4	52.4 %	± 1.06%	75 mm bearing ball
82	0.501	0.508	0.535	0.259	0.260	0.263	0.518	0.511	0.491	0.507	50.7			
83	0.537	0.549	0.582	0.292	0.297	0.306	0.544	0.542	0.526	0.537	53.7			
84	0.517	0.526	0.556	0.280	0.283	0.290	0.541	0.538	0.522	0.534	53.4			
85	0.529	0.540	0.572	0.287	0.291	0.300	0.542	0.539	0.523	0.535	53.5			
86	0.484	0.491	0.517	0.261	0.263	0.268	0.538	0.535	0.519	0.531	53.1			
87	0.508	0.516	0.543	0.273	0.275	0.280	0.536	0.533	0.517	0.529	52.9			
88	0.508	0.515	0.542	0.267	0.269	0.272	0.526	0.522	0.503	0.517	51.7			
89	0.523	0.534	0.566	0.272	0.275	0.281	0.521	0.515	0.496	0.511	51.1			
90	0.537	0.548	0.582	0.283	0.287	0.294	0.527	0.523	0.505	0.518	51.8			
91	0.563	0.577	0.615	0.513	0.532	0.563	0.911	0.923	0.916	0.917	91.7	91.5 %	± 1.42%	90mm Christmas ball
92	0.554	0.566	0.602	0.495	0.513	0.542	0.894	0.907	0.900	0.900	90.0			
93	0.553	0.567	0.603	0.509	0.525	0.552	0.920	0.927	0.915	0.920	92.0			
94	0.538	0.549	0.581	0.476	0.490	0.511	0.886	0.892	0.880	0.886	88.6			
95	0.545	0.557	0.590	0.504	0.522	0.548	0.926	0.937	0.929	0.931	93.1			
96	0.554	0.568	0.605	0.514	0.533	0.562	0.928	0.938	0.930	0.932	93.2			
97	0.537	0.548	0.581	0.491	0.505	0.529	0.913	0.921	0.910	0.915	91.5			
98	0.577	0.591	0.630	0.526	0.542	0.567	0.911	0.916	0.900	0.909	90.9			
99	0.563	0.577	0.614	0.516	0.531	0.556	0.916	0.921	0.905	0.914	91.4			
100	0.525	0.535	0.566	0.484	0.499	0.523	0.923	0.933	0.925	0.927	92.7			

B.3 Reflectivity measurements of randomly selected mirrored spheres

Measurement No. (N)	Paper patch			Probe patch			Calculated Reflectivity			averaged reflectivity $\bar{\rho}_{\text{sample (N)}}$		overall averaged reflectivity $\bar{\rho}_{\text{sphere}}$	standard deviation σ_{sphere}	Sphere type
	R	G	B	R	G	B	R	G	B	decimal value	%			
101	0.381	0.392	0.402	0.231	0.231	0.223	0.607	0.589	0.554	0.584	58.4	58.5 %	± 1.68%	100mm yard globe
102	0.388	0.400	0.409	0.241	0.242	0.234	0.622	0.605	0.571	0.599	59.9			
103	0.384	0.396	0.405	0.241	0.242	0.235	0.627	0.611	0.580	0.606	60.6			
104	0.388	0.400	0.410	0.228	0.227	0.219	0.587	0.568	0.534	0.563	56.3			
105	0.387	0.399	0.409	0.241	0.242	0.234	0.623	0.606	0.572	0.601	60.1			
106	0.374	0.386	0.395	0.236	0.237	0.230	0.630	0.613	0.582	0.608	60.8			
107	0.432	0.443	0.454	0.260	0.259	0.249	0.602	0.584	0.548	0.578	57.8			
108	0.378	0.390	0.399	0.228	0.227	0.219	0.603	0.583	0.548	0.578	57.8			
109	0.373	0.384	0.394	0.222	0.222	0.215	0.595	0.577	0.546	0.573	57.3			
110	0.402	0.414	0.423	0.238	0.236	0.226	0.591	0.571	0.535	0.565	56.5			
111	0.377	0.389	0.399	0.232	0.233	0.226	0.615	0.598	0.565	0.593	59.3	59.0 %	± 0.61%	150mm yard globe
112	0.338	0.349	0.358	0.206	0.206	0.200	0.609	0.591	0.558	0.586	58.6			
113	0.373	0.385	0.395	0.226	0.227	0.220	0.606	0.589	0.557	0.584	58.4			
114	0.337	0.348	0.357	0.205	0.205	0.199	0.608	0.590	0.557	0.585	58.5			
115	0.341	0.352	0.362	0.209	0.209	0.202	0.613	0.594	0.560	0.589	58.9			
116	0.328	0.339	0.349	0.205	0.205	0.199	0.623	0.604	0.571	0.599	59.9			
117	0.346	0.357	0.366	0.214	0.214	0.208	0.619	0.601	0.568	0.596	59.6			
118	0.358	0.370	0.380	0.217	0.218	0.211	0.606	0.588	0.556	0.584	58.4			
119	0.353	0.364	0.373	0.219	0.219	0.212	0.621	0.602	0.568	0.597	59.7			
120	0.337	0.348	0.358	0.205	0.205	0.199	0.608	0.589	0.555	0.584	58.4			

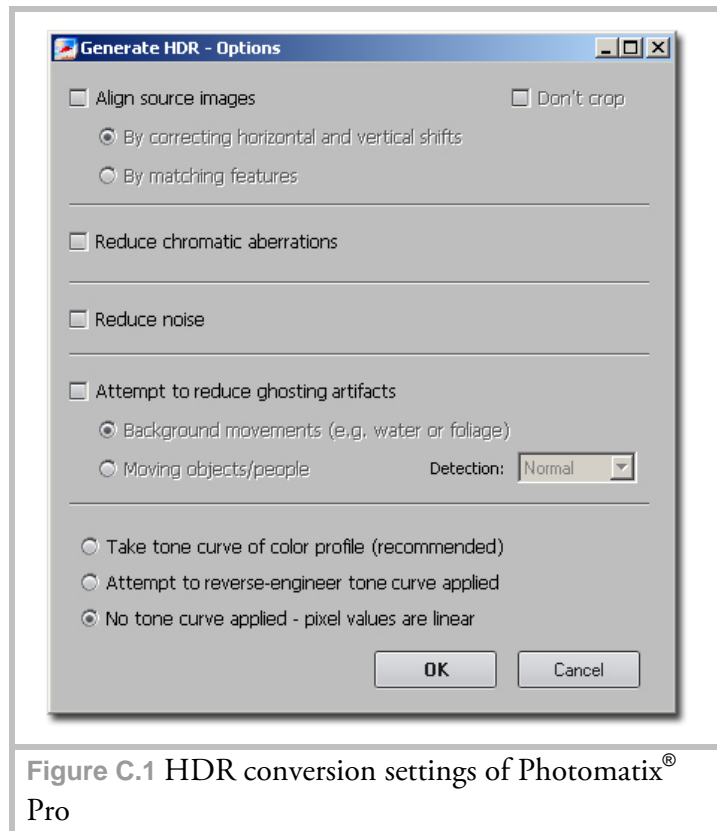
Appendix C

Settings for Photomatix® Pro

Although Photomatix® Pro offers a batch conversion tool (found in “Automate > Batch Processing...”), several tests have shown that high resolution linear TIFF images can not be processed correctly in that mode. Therefore manual processing of image sets to individual HDRs is recommended. The required settings are shown in figure to the right. We assume linear TIFF images (converted from RAW files) to be the input data for HDR conversion.

The listed options are selected as follows (in the order of their appearance):

- image alignment is recommended only if few images of up to 9 megapixels resolution (reference value determined during testing) are to be combined. At least 4 GB of RAM are required for this operation. Often, processing large files will crash the operating system;
- chromatic aberrations reduction may help to fix color spill problems around the edges of an image taken by a fisheye lens. For most IBL applications this option is not required;
- noise reduction may help if an image sequence was shot in low-light conditions, although the HDR process itself (having a sufficient number of differently exposed images) may significantly reduce noise levels;
- moving objects in the scene produce noticeable ghosting artifacts which the software may try to reduce. Not recommended for IBL applications;
- as the pixel values of the converted TIFF images are linear the corresponding option is chosen.



The converted images are saved as 32-bit OpenEXR images with PIZ-Compression.

Appendix D

Settings for Autopano[®] Pro

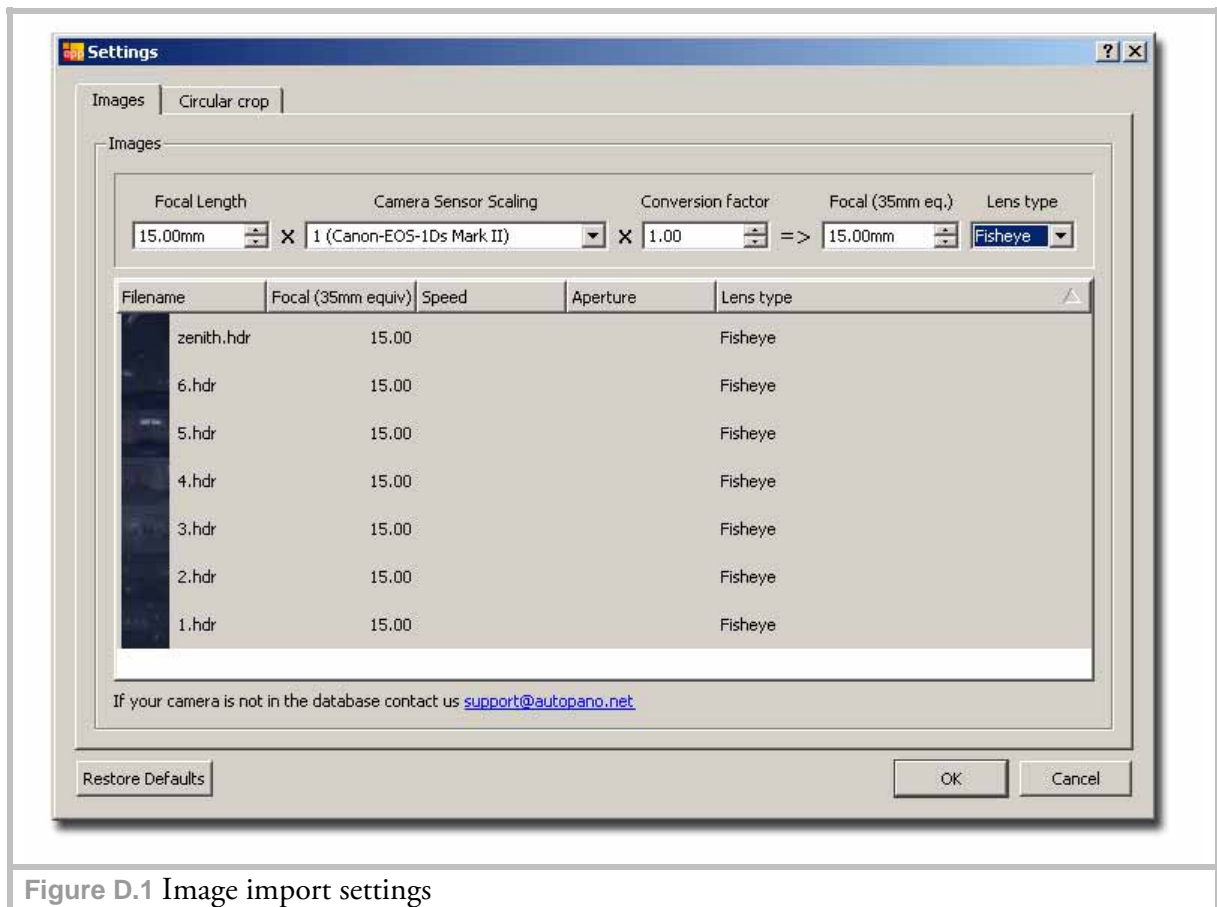


Figure D.1 Image import settings

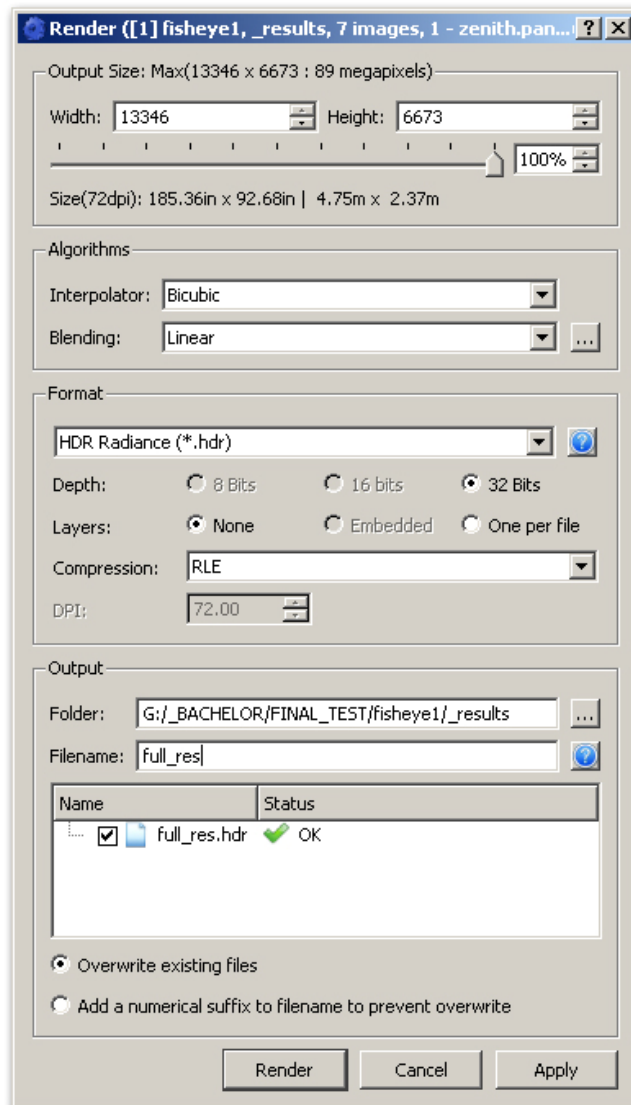
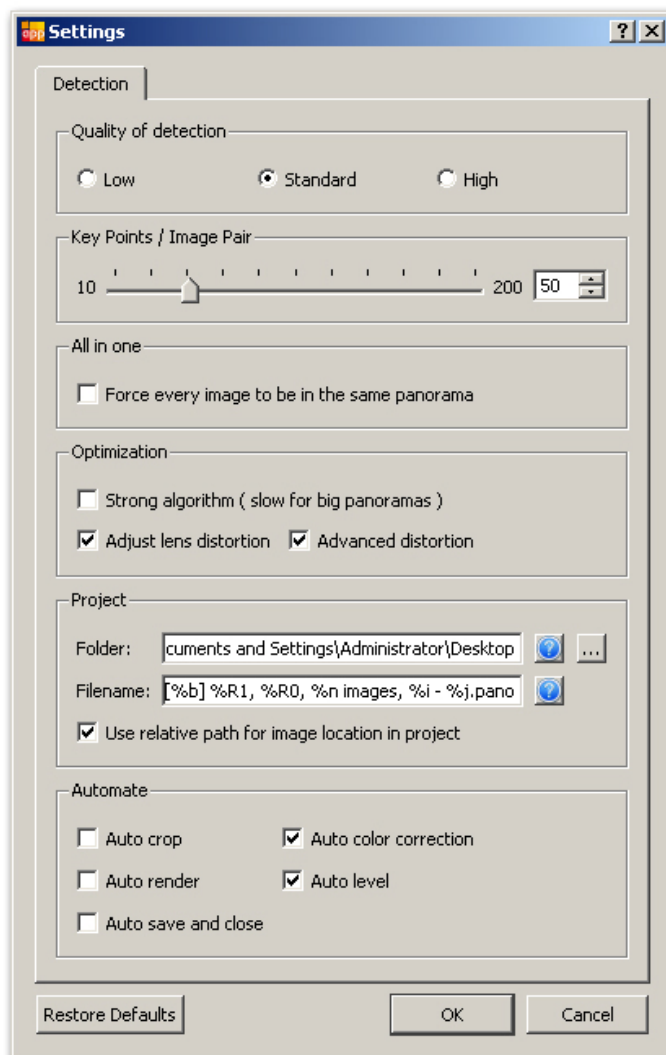


Figure D.2 Detection settings (left) / Panorama output settings (right)

Notes on Autopano® Pro settings

It is important to leave “Auto Level” and “Auto Color” options in the “Automate” settings group turned on, as they allow the software to temporarily apply adjustments to the HDR image set so that this becomes optimized for automatic panoramic image alignment.

Although Autopano® Pro offers different advanced image blending algorithms (like Multiband and Smartblend), these algorithms may produce insufficient results when overlapping light sources are to be blended in two or more views. Effects like dark halos may appear and thus it is best to use the linear blending mode instead. Several tests using this mode produced expected results.


The interpolation method affects the overall sharpness of the final image. Though all test images were rendered using the “bicubic” interpolation, other methods may produce better results at the expense of rendering time.

Rendering results



- Rendering results 1-

Rendering results



Mirrored ball

- Rendering results 2-

Rendering results

