





Tutorial Luminance Maps for Daylighting Studies from High Dynamic Range Photography

C. Pierson ^a, C. Cauwerts ^a, M. Bodart ^b, and J. Wienold ^c

^aArchitecture et Climat, Faculté d architecture, d ingénierie architecturale, d urbanisme (LOCI), Université catholique de Louvain (UCLouvain), Louvain-la-Neuve, Belgium; ^bFaculté d architecture, d ingénierie architecturale, d urbanisme (LOCI), Université catholique de Louvain (UCLouvain), Louvain-la-Neuve, Belgium; ^cLaboratory of Integrated Performance in Design (LIPID), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

BSTR CT

In the field of lighting, luminance maps are often used to evaluate point-in-time lighting scenes from the occupant's vantage point. High Dynamic Range (HDR) photography can be used to generate such luminance maps. The aim of this tutorial is to present a comprehensive overview of a step-by-step procedure to generate a 180° luminance map of a daylight scene from a sequence of multiple exposures with semiprofessional equipment and the Radiance suite of programs. The procedure consists in capturing a sequence of multiple exposures of the visual scene; selecting the useful exposures; merging the exposures to generate the HDR image by using the predefined camera response function; nullifying the exposure value; resizing and cropping the HDR image by using the predefined fisheye view coordinates; adjusting the projection of the HDR image by using the predefined distortion function; correcting the vignetting of the HDR image by using the predefined vignetting curves; correcting the alterations of the HDR image due to the Neutral Density (ND) filter if one was used, by using the predefined ND correction function; adjusting the photometry of the HDR image by using the measured spot luminance value; editing the HDR image header by using the predefined projection type and real viewing angle; and checking the validity of the HDR image by using the measured vertical illuminance, and, if needed, the predefined luminous range. To conclude, an analysis of errors is made and attention points to adapt the procedure for electric or circadian lighting studies are discussed.

ARTICLE HISTORY

Received 2 April 2019
Revised 21 October 2019
Accepted 21 October 2019



KEYWORDS

High Dynamic Range (HDR); luminance map; calibration; daylighting; photography; light measurement

Learning outcomes

It is our intent that, after studying this article, the reader will:

- (1) be able to explain the concept of HDR photography – why is it used, and what is the basic concept;
- (2) feel familiar with some basic concepts related to digital imaging and photography that are required to understand the calibration procedure: *no-parallax point*, *vignetting*, *camera response function*, *pixel overflow*, *lens flare*, *projection types*, *bracketing*, *image header*;
- (3) have an overview of the multiple steps required to create (i.e. generate and calibrate) a luminance map using affordable equipment and free software, and understand the reasons why each of these steps is necessary;
- (4) be able to generate an HDR image of a daylight visual scene, and calibrate it radiometrically, photometrically, and geometrically;
- (5) understand what is a reasonable range for each calibration feature, and for the luminance and vertical illuminance errors of a calibrated HDR image;
- (6) be aware of the limitations and pitfalls associated with the use of HDR photography for luminance measurement in daylighting, as well as the challenges when adapting the technique to electric lighting and color studies.

CONTACT C. Pierson  clotilde.pierson@uclouvain.be  Architecture et Climat, Faculté d architecture, d ingénierie architecturale et d urbanisme (LOCI), Université catholique de Louvain (UCLouvain), Place du Levant, 1/L5.05.04 B-1348, Louvain-la-Neuve, Belgium.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/ulks.

© 2019 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

1 Introduction

In the field of lighting, luminance maps are often used to evaluate point-in-time lighting scenes from the occupant's vantage point. High Dynamic Range (HDR) photography can be used to generate such luminance maps. HDR photography consists of capturing multiple exposures of a scene from an identical vantage point and merging them in order to obtain an image with a higher range of brightness, namely an HDR image. In traditional images, also called Low Dynamic Range (LDR) images, the contrast ratio is limited by the pixel range and the response function of the camera. For instance, the camera used for the examples given in this tutorial can capture LDR images with a maximal contrast ratio of around 33:1. Though indoor daylit scenes can have a dynamic range up to 10^9 :1 if the sun is in the field of view (FOV) (Jakubiec et al. 2016b). After calibration, a luminance map is produced from the HDR image, and can be used to derive much useful information such as the average luminance and glare source(s) in the FOV, or the vertical illuminance.

Nowadays, there exist two main methods of generating calibrated HDR images: absolute calibration and step-by-step calibration with an automatic merging algorithm. The absolute calibration consists of finding, for each exposure, a direct relationship between each pixel value and the luminance of the point it represents in the scene, with the relationship being applicable for all luminance levels that can be found in the field. From these relationships, the luminance validity range of each exposure can be defined, and the merging of multiple exposures into an HDR image is done accordingly (Dumortier et al. 2005). This method requires proper equipment and coding skills; hence it is principally used by professional manufacturers to produce luminance cameras. Due to the demanding calibration process required for high-precision solutions, luminance cameras tend to be expensive and often not commercially viable for lighting research (Wüller and Gabele 2007). The step-by-step calibration consists of applying an automatic algorithm that merges multiple exposures into an HDR image while doing a radiometric calibration, before manually adjusting the generated HDR image to calibrate it photometrically and geometrically. There exist several automatic merging algorithms such as *hdrgen* from *Radiance* (Ward 1998) or

pfscalibration from *pfstools* (Mantiuk et al. 2007), although they are generally based on the same principle (see 2.4.3). This step-by-step calibration makes it possible to produce luminance maps with good photometric accuracy, but without requiring expensive equipment nor specific coding skill. Several studies (Cai and Chung 2011; Inanici 2006; Stanley 2016) have shown that, for scenes where the maximal luminance value was lower than $30,000 \text{ cd/m}^2$, most luminance values from the luminance map could be expected within a 10% error range, although higher errors also occurred. This error margin makes HDR photography relevant as a measurement tool, but the step-by-step calibration is tedious and error-prone since it involves multiple steps.

The aim of this tutorial article is to present a synthetic overview of the step-by-step calibration process required to generate an HDR image that can be used to study the luminance distribution in an interior daylit scene, and to propose a practical procedure to execute it. The scope of this tutorial is restricted to the generation of luminance maps covering a 180° FOV from a sequence of jpeg images, and requires therefore the use of a 180° fisheye lens. However, the steps described in the tutorial can be adapted for the capture of a sequence of raw images or the use of a normal (non-fisheye) lens, on condition that this lens is not a zoom lens. Unless the zoom is fixed, it is not possible to ensure that the lens is in the exact same position during the entire calibration procedure and for all following captures of LDR sequences, which causes a reproducibility issue.

A thorough description of each step of the process is provided, as well as requirements, options, and common mistakes. Examples performed using the equipment described below are also presented in boxes at the end of some sections to give an idea of what should be expected as calibration features. References are made to useful literature, and attention points are given to adapt the method for electric or circadian lighting studies.

2 Tutorial

2.1. Requirements

2.1.1. Software

This tutorial to generate calibrated HDR fisheye images is based on several freely accessible software,

namely *qDslrDashboard* (Hubai 2014) (available for Windows, Linux, and macOS) to remotely control the camera; the *Radiance* suite of programs (Ward 1987) (available for Windows, Linux, and macOS) and the command tool *hdrgen* (Ward 1998) (available for Linux and macOS) to process the images; and *Photosphere* (Ward 1998) (available for macOS) to apply the photometric adjustment. It is recommended to use the latest official release of these software packages. Links to the software packages are given in the references. Novice users are invited to consult tutorials available on Radiance website (<https://www.radiance-online.org/learning/tutorials/>) to get familiar with *Radiance* tools, as well as the discussions available on Radiance and HDRI forums (<https://discourse.radiance-online.org/>) for further help.

For the examples given in this tutorial, version 3.5.5 of *qDslrDashboard* was installed on a Windows 10 computer, although a new version (3.5.7) has been released since then. Version 5.2 of *Radiance*, which includes the programs *falsecolor*, *ra_xyze*, *pcomb*, *pcompos*, *pfilt*, *pvalue*, *ximage*, *getinfo*, and *evalglare*, as well as *hdrgen* were used on Linux Ubuntu 16.04 LTS, which was virtualized through version 12.5.5 of VMWare Workstation on a Windows 10 computer. The latest version of *evalglare* (2.08) was used, although this version was not yet updated in *Radiance* 5.2. Version 1.8.21U of *Photosphere* was also used on macOS Sierra, although a new version (1.8.25U) has been released in the meantime.

2.1.2. Hardware

The following hardware equipment is required to generate calibrated HDR fisheye images in order to analyze daylit scenes. This equipment is needed each time a new calibrated HDR fisheye image has to be generated:

- a Digital Single-Lens Reflex (DSLR) camera (CCD and CMOS sensors are equivalent competitors in the camera market nowadays (Stanley 2016))
- a circular fisheye lens which is compatible with the camera (fisheye lenses are frequently

used in lighting research to maximize the captured FOV)

- a tripod on which to set the camera to ensure that the camera does not move during the sequence of multiple exposures
- a computer to remotely control the camera and process the images
- a calibrated spot luminance meter and a middle grey target to make the photometric adjustment through a spot luminance measure on the target (L_T). In case of expected very high luminance (e.g. sun seen through fabric), it is useful to use a translucent target, which should be lit by the sun and placed on the window. Both grey and translucent targets should show near to lambertian behavior to limit the influence of changing luminance values under different viewing angles.
- a calibrated illuminance meter to check the validity of the calibrated HDR fisheye image
- an ND filter could be needed if the sun or a very bright surface is in the FOV

Some of the equipment, though, is only needed to conduct the one-time setup for a specific camera/lens association (see 2.3):

- a color chart
- a panoramic rotation unit and a sliding plate to move the camera on the tripod
- a stable and bright light source (e.g. a halogen spotlight) in a dark room, or a stable overcast sky simulator (e.g. a mirror box) with a semicircular platform

Before starting the step-by-step calibration, the proper functioning of the equipment should be ensured, especially regarding the camera and the lens.

Because it deteriorates over time, the camera sensor should be checked for stuck, dead, or hot pixels (Jacobs and Wilson 2007). Stuck pixels are damaged pixels that will always appear colored or overexposed, even if no light reaches the sensor. Therefore, by taking a picture with the lens cap on the camera, stuck pixels are detected as the colored or white spots in the image. Dead pixels are permanently damaged pixels that do not receive any power. They will thus always appear black, even when the image is fully overexposed. Hot pixels are not permanently damaged, but show up during long exposures as the

sensor heats up. To detect hot pixels, a picture should be taken with the lens cap on using a long exposure. Every non-black and non-stuck pixel in that long-exposed image is a hot pixel. Hot pixels should not be a problem since for daylight scenes, the maximal exposure time should be kept under 2 seconds to prevent the lighting conditions to vary too much. However, the presence of several stuck or dead pixels requires to change camera to avoid distorted measurements.

The reproducibility of the camera should be checked as well. This can be done easily in a stable indoor scene, under a non-flickering (see 3.2) light source. The camera is set on the tripod and remotely controlled by the computer to ensure immobility. An LDR image is captured for each pertinent lens aperture with all other settings fixed. Without changing the situation nor moving the camera, a second series of LDR images is captured for each previous aperture and the exact same settings. In an image editing software, the LDR images can be analyzed pixel by pixel. If non-black corresponding pixels in two LDR images of same aperture have the same RGB values, the images captured by the camera are reproducible. This test could also be conducted by varying the exposure time instead of the lens aperture.

For the examples given in this tutorial, the equipment that was used consists of a Canon EOS 5D Mark II camera and a Sigma EX DG 8mm f/3.5 fisheye lens. The camera is remotely controlled by a computer with *qDslrDashboard* via USB, and the images are processed with *Radiance* tools, *hdrgen*, and *Photosphere*. The ND filter used in specific cases is a Kodak Wratten 3.0 ND filter. Furthermore, a Hagner digital luxmeter EC1-X and a Konica Minolta LS-110 luminance meter, both calibrated in the 12 months preceding the data acquisition, are used. The sensor of the camera was checked with the described methods and seems to be clear of stuck and dead pixels, but shows several hot pixels for an exposure time of 30 seconds. The camera also provides reproducible results.

2.1.3. Camera settings

The camera settings should be set according to [Table 1](#), as it has been recommended in the literature

Table 1 Camera settings that should be used to capture a sequence of multiple exposures for HDR photography.

Setting	Value
Film speed	ISO 100
White balance	Daylight (5200K)
Exposure mode	Manual
Light metering mode	Insignificant
Focus mode	Manual
Focus value	Infinite
Image quality	Largest
Image type	JPEG or RAW
Picture style	Neutral
Peripheral illumination correction	Disabled
Color space	sRGB

(Inanici 2009; Reinhard et al. 2006). These settings have to be kept constant for the whole calibration process and every luminance maps generated afterwards. For instance, modifying the white balance setting during the process would alter the color space transitions (Inanici 2006) and the camera response function (Jacobs 2012), which would result in calibration errors. Also, the focus value, if not locked, could vary between each exposure, which would hinder the merging of these exposures.

Some settings from [Table 1](#) might be referenced differently when using another camera brand. However, it is important to make sure that all image effects or corrections applied by the camera are disabled or set to neutral. Regarding aperture and mean exposure time, these settings will be determined below (see 2.4.1) and are set through the remote control software.

2.2. Procedure

The complete procedure to create a 180° luminance map from a visual scene with daylight is shown in [Fig. 1](#). Each step of this procedure is explained in detail in this article (2.4 and 2.5), and the lines of code to execute the steps in *Radiance* are given. The procedure consists in capturing a sequence of multiple exposures of the visual scene; selecting the useful exposures; merging the exposures to generate the HDR image by using the predefined camera response function; nullifying the exposure value; in resizing and cropping the HDR image by using the predefined fisheye view coordinates; adjusting the projection of the HDR image by using the predefined distortion function; correcting the vignetting of the HDR image by using the predefined

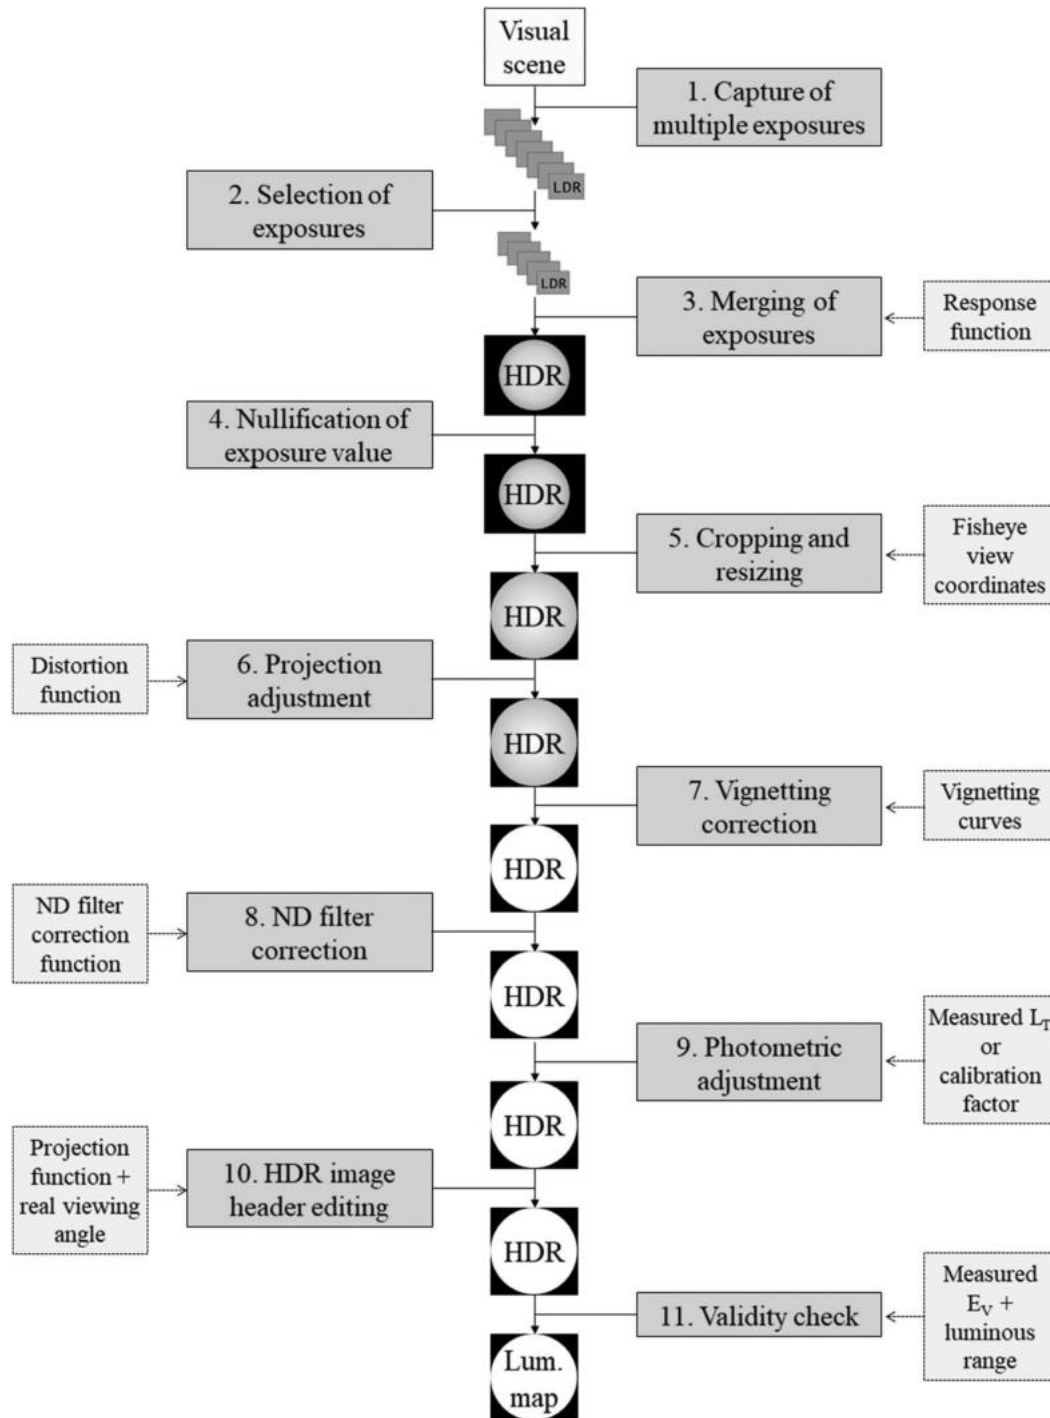


Fig 1 Step-by-step procedure to create a 180° luminance map from a visual scene with daylight using HDR photography.

vignetting curves; correcting the alterations of the HDR image due to the ND filter if one was used, by using the predefined ND correction function; adjusting the photometry of the HDR image by using the measured spot luminance value; editing

the HDR image header by using the predefined projection type and real viewing angle; and checking the validity of the HDR image by using the measured vertical illuminance, and if needed the predefined luminous range.

Some calibration material, such as the camera response function or the vignetting curves, has to be defined prior to executing the procedure. The methods used to derive such calibration material are explained in the one-time setup (2.3). These steps should be performed only once per equipment, and can be used afterward for each luminance map of a daylit scene generated with the same equipment. Therefore, the specific equipment required for the one-time setup (a color chart, a panoramic rotation unit and a sliding plate to move the camera on the tripod, and a stable and bright light source in a dark room or a stable overcast sky simulator with a semi-circular platform) should only be available once. The rest of the equipment will always be needed for the HDR image generation and the HDR image calibration, except for the ND filter, of which the use depends on the lighting conditions.

2.3. One time setup

2.3.1. Determination of the no-parallax point real viewing angle and fisheye view coordinates

The no-parallax point or entrance pupil point, sometimes wrongly referred to as the nodal point (Kerr 2008), is the point of the lens at which the rays entering the lens should converge. For traditional lenses, this point can be determined using a tripod with a panoramic rotation unit and a sliding plate on which the camera is set, as well as two vertical markers aligned with the centre of the rotation unit, in front of the camera (Jacobs 2012). If the centre of rotation of the camera coincides with the no-parallax point, the rear marker stays hidden by the one in front as the camera rotates (Fig. 2).

In photography, knowing the position of the no-parallax point is important when panoramic images are captured by rotating the camera for instance. For fisheye lenses, the no-parallax point is unfortunately not unique, and slides along the lens optical axis, closer to the front of the lens as the angle of view increases (Anderson 2006). There is therefore a range of least-parallax points. For the following steps, one of the least-parallax points (i.e. the one in the middle of the range) is selected, and the hypothesis is made that this point is the lens no-parallax point. The camera has to be rotated around this hypothesized no-parallax point in the future to avoid parallax error,

namely a difference in the apparent relative position of objects due to a displacement of the line of sight.

Besides the fact that the edge of a fisheye view is in general blurry, the real viewing angle of a fisheye lens might not be exactly 180° and the centre of a fisheye view might not be exactly at the centre of the image (Jacobs 2012). To determine the total viewing angle of a fisheye lens, the no-parallax point has to be positioned between two reference grids. The no-parallax point should be aligned with the middle vertical lines of both grids, and centred at equal distance from these grids, so that the spacing between two lines of the grids corresponds to one degree in the camera FOV (Fig. 3).

When the camera is looking at 0° , the two middle vertical lines of the grids should appear at the border of the fisheye LDR image, as well as when the camera is looking at 180° . In an image editing software, the distance in pixels between these lines in the 0° and 180° LDR images can be evaluated. This distance should be the same in both images. If additional vertical lines – over the 90° middle lines – appear in both images, the total viewing angle is thus larger than 180° and can be defined according to the number of additional lines visible.

Furthermore, by using any LDR image captured with the camera and lens, the fisheye view coordinates can be defined through an image editing software. The fisheye view coordinates are the coordinates (in pixels) of the bottom-left corner of the square encompassing the total viewing angle, as well as the fisheye view diameter. It should be noted that the origin of the coordinates system might be different between the image editing software (e.g. the origin is the upper left corner in *Gimp*, and in *Radiance*, the origin is typically the bottom left corner).

The no-parallax point of the example equipment is located 2 mm in front of the lens gold ring, the real viewing angle is 186° , and the center of the fisheye view is slightly shifted towards the upper-right corner of the image.

2.3.2. Determination of the luminous range

The equipment luminous range is characterized by the minimal and maximal luminance values that the camera/lens combination can measure through HDR photography (Jakubiec et al. 2016a). In an LDR

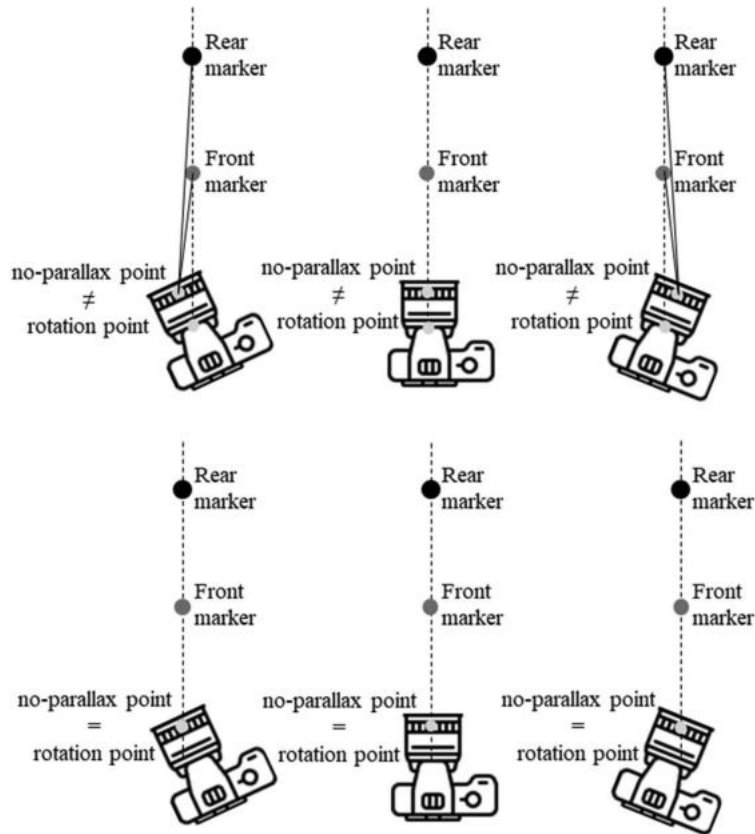


Fig 2 Setup of the equipment to evaluate the no-parallax point.

image, these extreme values are function of the sensitivity and the resolution of the camera sensor, of the lens, and of the ISO, aperture and exposure time settings. Since in HDR photography, the exposure time is varied, and a low ISO is recommended to minimize the noise in the image, the minimal and

maximal luminance values measured in an HDR image depends on the camera sensor and lens, and on the aperture setting. Therefore, for a specific camera/lens combination, the luminous range can be defined for each aperture.

The minimal luminance value, namely the lower limit of the luminous range, should theoretically be zero. Due to noise and potential stuck pixels however, a slightly higher value should be expected. For typical indoor daylight scenes though, the noise floor at the lower end of the luminous range should not matter much.

To determine the maximal luminance value that can be captured for one aperture, a sequence of multiple exposures (see 2.4.1) of a very bright light source such as the sun has to be captured in the centre of the fisheye view to avoid the vignetting effect. Exposures are then selected (see 2.4.2), and an HDR image is generated from the sequence (see 2.4.3). The maximal luminance value is defined as




Fig 3 Setup of the equipment to evaluate the real viewing angle and the centre of the fisheye view.




Fig 4 Maximal luminance values that can be captured through HDR photography with the example equipment for different apertures.

the peak (maximal) luminance in this HDR image, which can be retrieved using *falsecolor*. As the aperture size increases, the maximal luminance value decreases. A larger aperture size will indeed allow more light to fall on the sensor in a given period of time. This will cause more images in a sequence of multiple exposures to be overexposed, and thus to display a reduced contrast between the bright light source and the background. The dynamic range of the final HDR image will then be lower for a large aperture size than for a small aperture size.

Attention should be paid to the fact that *Radiance* uses the equal-energy white point [(x, y, z) = (0.33, 0.33, 0.33)] for its calculations, and slightly different RGB primaries than sRGB primaries. Consequently, a luminance outputted by *Radiance*-based programs, such as *ximage*, *falsecolor* and *evalglare*, is calculated according to Equation 1.

$$L = 179 * (R * 0.2651 + G * 0.6701 + B * 0.0648) \quad (1)$$

where L is the luminance of the pixel [cd/m^2];

179 is the standard luminous efficacy [lm/W] in *Radiance*;

R , G , and B are the spectrally weighted radiance of the pixel [$\text{W}/\text{m}^2\text{sr}$];

0.2651, 0.6701, 0.0648 are calculated from CIE chromaticities for the reference primaries (*Radiance* primaries) and the equal-energy white point as used by *Radiance*.

Most HDR image merging programs, such as *hdrgen* and *Photosphere*, use sRGB as their default output color space. sRGB white point is the one from CIE Standard Illuminant D65 [(x, y, z) = (0.3127, 0.3290, 0.3583)], thus slightly different from the equal-energy white point. Consequently, luminance calculation from HDR images should be performed using Equation 2, as it is executed in *Photosphere* (Inanici 2006).

$$L = 179 * (R * 0.2126 + G * 0.7152 + B * 0.0722) \quad (2)$$

where L is the luminance of the pixel [cd/m^2];

179 is the standard luminous efficacy [lm/W] in *Radiance*;

R , G , and B are the spectrally weighted radiance of the pixel [$\text{W}/\text{m}^2\text{sr}$];

0.2126, 0.7152, 0.0722 are calculated from CIE chromaticities for the reference primaries (sRGB) and the CIE Standard Illuminant D65 white point.

Although the difference between the two equations looks minor, the error in terms of luminance for these two equations was evaluated with the example equipment under LED lighting (color temperature of 5400K) using Equation 3.

$$L = \frac{L_{\text{HDR}} - L_{\text{meas}}}{L_{\text{meas}}} * 100 \quad (3)$$

where L is the error in luminance [%];

L_{HDR} is the luminance derived from the HDR image using equation 1 or 2 [cd/m^2];

L_{meas} is the luminance measured in the scene with a spectroradiometer [cd/m^2].

In average, the error between the measured and HDR-derived luminance for 18 colored samples from a Macbeth chart was observed to be around two times larger for the luminance derived from the HDR image using Equation 1 (8.8%) than for the luminance derived from the HDR image using Equation 2 (4.6%). No luminance error difference was observed between the two equations for neutral (grey) samples (both 2.1%). Nevertheless, using *ximage* and *falsecolor* to retrieve luminance value is acceptable in a context of indoor scenes. Indeed, the visual environment inside buildings is relatively neutral, making the use of Equation 1 acceptable. The line of code in *Radiance* to print the peak (min. and max.) luminance values of an HDR image (*HDRfile.hdr*) with *falsecolor* in a false-colored HDR image (*HDRfile_peak.hdr*) is as follows:

```
> falsecolor -i HDRfile.hdr -e > HDRfile_peak.hdr
```

If the peak luminance in an HDR image generated with a specific equipment is near the upper limit of the equipment luminous range, the picture might exhibit pixel overflow. A solution to avoid pixel overflow is to use an ND filter (see 2.3.6).

The upper limit of the luminous range for the example equipment was evaluated for each useful aperture by capturing a sequence of 15 images (exposure time going from 1/8000 to 2 seconds) of the sun under clear sky conditions. These upper limits, shown in Fig. 4, are similar to the ones presented by Jakubiec et al (Jakubiec et al. 2016b).

2.3.3. Determination of the response function

The camera response function is required to generate an HDR image when the LDR images are in the jpeg format, but not when they are in the raw format (see 2.4.1). The response function relates pixel values to incoming light quantities (Hirning 2014). More specifically, it matches radiance values from the scene ($\text{W}/\text{m}^2\text{sr}$) to pixel values (0–255) for the RGB channels. Each camera has its specific response function, even cameras from the same model (Jacobs 2007).

To approximate the camera response function, several algorithms have been developed, amongst which the one from Debevec and Malik (Debevec and Malik 1997) which was later modified by Mitsunaga and Nayar (Mitsunaga and Nayar 1999), and the one from Robertson et al (Robertson et al. 2003). These algorithms rely on a series of aligned exposures. By looking at the evolution of one pixel value across the exposures, one can sample the shape of the camera response function for the corresponding part of the curve. Using linear optimization (Debevec and Malik 1997) or a polynomial approximation (Mitsunaga and Nayar 1999) with several pixels, a smooth response function can be obtained for each channel (Reinhard et al. 2006). *hdrgen* and its implementation in *Photosphere* approximate the RGB camera response curves using Mitsunaga and Nayar’s method, whereas *pfscalibration* offers the option to choose between Robertson’s or Mitsunaga and Nayar’s algorithms. These algorithms all require a sequence of LDR images taken with a specific camera as input to derive the response function of that camera.

Since the response function depends on the white balance setting of the camera and the principal lighting source (Jacobs 2012), it should be derived once in correct conditions, and reused for all HDR images generated from the same camera with the same white balance setting (“daylight – 5200K” for daylight scenes) and under a similar main lighting source (daylight, in this case). The scene photographed for the sequence of multiple exposures has to be chosen carefully to optimize the response recovery. In the case of daylight, it is recommended to select an interior daylit scene with large and smooth gradients throughout interior and exterior views, with very bright and dark areas, and as non-chromatic as possible (Fig. 5) (Inanici 2006; Inanici and Galvin 2004; Jacobs 2012; Reinhard et al. 2006). Furthermore, large apertures of the lens (e.g. f/3.5 or f/5.6) increase the probability of biasing the response curve derivation due to the vignetting effect on the LDR images of the sequence (see 2.3.5). To limit this effect, a small aperture such as f/16 should be preferred. At last, the indications for the capture of multiple exposures (see 2.4.1) should be taken into consideration for the response function derivation.




Fig 5 Example of a visual scene appropriate for the derivation of the camera response function.

Once a sequence of LDR images is captured, the RGB response curves of the camera are evaluated by *hdrgen* on the basis of this sequence. The line of code to generate the camera response function (*response_function.rsp*) from a sequence of multiple exposures (*LDRfiles_*.JPG*) with *hdrgen* is as follows:

```
> hdrgen LDRfiles_*.JPG -o HDRfile.hdr -r response_function.rsp -a -e -f -g
```

The algorithm outputs approximated response curves in an *.rsp* text file containing the polynomial order of the three RGB curves and their coefficients. Attention should be paid to the direction in which the coefficients are printed since it is

reversed between *hdrgen* (from k_N to k_0) and *Photosphere* (from k_0 to k_N) for instance.

Regarding the example equipment, three response functions were derived from three scenes that matched the required characteristics, similar as the scene in Fig. 5. RGB response curves were derived for the three scenes and compared. They showed very similar behavior (Fig. 6). Final compound camera response curves for the example equipment were derived from the means of the three curves. An *.rsp* text file was created that contained the polynomial order of the mean RGB curves and their coefficients in the order of *hdrgen*.

2.3.4. Determination of the projection and distortion functions

Fisheye lenses project 3D scenes on 2D images with a specific projection type (Fig. 7) amongst equidistant (i.e. equiangular), equisolid-angle (i.e. equal-area), orthographic (i.e. hemispherical), and stereographic (i.e. planispheric) (Cauwerts et al. 2013).

The two most widely implemented projection types in commercially available fisheye lenses nowadays are the equidistant and equisolid-angle ones. Although each lens model is related to one theoretical projection, the implementation of the projection is not perfect in practice and suffers from small inaccuracies. For several applications, the geometrical accuracy of an HDR image is primordial, hence the importance to know the real projection function of the lens and to be able to adjust it if need be. For instance, the calculation from an HDR image of any discomfort glare metric using *evalglare* (Wienold and Christoffersen 2006) relies on two geometrical quantities, i.e. the solid angle and the position index




Fig 6 RGB response curves derived with *hdrgen* for three different visual scenes captured with the example equipment and mean RGB response curves.

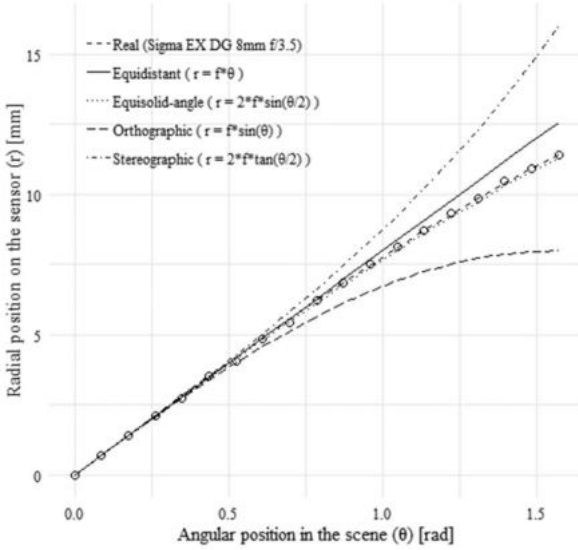


Fig 7 Theoretical projection types to project a 3D scene on a 2D image and measured projection of example equipment.

of the glare source(s). But the solid angle and the position index vary according to the projection function of the lens. *Evalglare* is currently able to devise these geometrical quantities for fisheye lenses using equidistant and orthographic projection types (Wienold 2012). Consequently, there is a need to adjust the projection of an HDR image if the lens that was used offers another projection type than those supported by *evalglare*. For applications where geometry of the visual scene is relevant, such as glare studies, it is generally recommended to have an HDR image with an equidistant projection.

Before adjusting the projection of an HDR image, the first step is to define the real projection function of the lens. The information about the theoretical projection type of the lens might be given by the manufacturer. Otherwise, one relatively easy way to verify this information is to measure it manually.

A first option to measure it is to place the camera facing a 90° corner of a room, with its hypothesized no-parallax point at equal distance between the two walls (Fig. 8). A scale is measured and printed on the walls, so that there is a reference in the scene marking each 5° angle of the 180° FOV, and a picture is captured in this setup. The radial position (in pixels) of the 5° angle references are extracted from the LDR image using an image editing software, and translated to the absolute radial position (in mm) on the camera sensor.

A second option to measure manually the projection type of the lens requires a panoramic rotation unit. First, an LDR image of a vertical marker on the horizon line of the fisheye view has to be taken every 5° by rotating the camera on its hypothesized no-parallax point until the lens FOV is covered. The radial position (in pixels) of this vertical marker can be easily extracted from each picture using an image editing software, and translated to the absolute radial position (in mm) on the camera sensor.

For both methods, each measured radial position is related to a certain angle of the lens, and can be displayed on a graph to be compared with the theoretical projection types (Fig. 7). By computing the Mean Squared Error (MSE) of each theoretical projection function against the measured values, the theoretical projection type of the lens can be defined as the one producing the smallest MSE. Once the theoretical projection type is known, the real projection function of the lens can be derived from the measured values with a regression based on the theoretical projection model, as shown in Equation 4 for a theoretical equisolid projection.

$$r = a^*f^* \sin \frac{\theta}{b} \quad (4)$$

where r is the distance on the sensor from the optical axis [mm];

f is the focal length of the lens [mm];

θ is the entrance angle measured from the optical axis [rad];

a and b are the parameters of the real projection function to be defined.

There exist other methods to derive the real projection function of a lens:

the method for which the no-parallax point of the camera/lens is set at the centre of a semi-circular platform, on which references are positioned every 5° in the 180° FOV of the camera (Cauwerts et al. 2013). By retrieving the pixel coordinates of each 5° reference of the platform from an LDR image taken by the camera, the real projection function of the lens can be derived.

the method for which the camera/lens is set in front of a planar reference grid with its no-




Fig 8 Setup of the equipment to evaluate the real projection function.

parallax point aligned with the centre of the grid at a known distance from it (Geisler-Moroder et al. 2016). By evaluating the distances between each grid line and the middle one, and comparing these distances with the pixel distances retrieved from an LDR image of the grid, the real projection function of the lens can be derived.

If the real projection function differs from the two most common theoretical projection functions (i.e. equidistant or orthographic) and if the application of the HDR image requires geometrical accuracy, a distortion function should be derived and applied to the HDR image to adjust its projection. The distortion function establishes the displacement that should be applied to a pixel according to its original position to get an equidistant projection in an HDR image.

Since the equisolid-angle type is one of the two most widely available lens types, the distortion function to adjust an HDR image with an equisolid-angle projection to an HDR image with an equidistant projection has already been defined (Ward 2016), as shown in Equation 5.

$$r_{equidistant} = \frac{2}{\pi} * \sin^{-1} \left(\sqrt{2} * r_{equisolid-angle} \right) \quad (5)$$

where $r_{equidistant}$ is the normalized distance from the optical axis for an equidistant projection [0 – 0.5];

$r_{equisolid-angle}$ is the normalized distance from the optical axis for an equisolid-angle projection [0 – 0.5].

For fisheye lenses having another projection type than the equisolid-angle one, the distortion

function can be approximated by a second-order polynomial without intercept, as shown in Equation 6. The approximation is done on the basis of the adjusted radial position of the pixels and their original radial position, both normalized in the [0– 0.5] range for the distortion function to be easily applied afterward. The adjusted radial positions can be derived using the theoretical equidistant projection function, whereas the original radial positions are either derived from the real projection function previously defined, or from the measured radial position of the vertical marker in the series of LDR images using the panoramic rotation unit for instance.

$$r_{equidistant} = a * r_{original}^2 + b * r_{original} \quad (6)$$

where $r_{equidistant}$ is the normalized distance from the optical axis for an equidistant projection [0 – 0.5];

$r_{original}$ is the normalized distance from the optical axis for the real projection of the lens [0 – 0.5];

a and b are the parameters of the distortion function to be defined.

At last, the distortion function has to be applied on the pixels of the HDR image to adjust. For this purpose, the distortion function is stored in the `fish-eye_corr.cal` file (Ward 2016) in place of the `mapsolid r` equation. The default distortion function is the equisolid-to-equidistant function, but it can be changed if need be. Besides adjusting the projection of an HDR image, the application of the `fish-eye_corr.cal` file also sets the pixels of the HDR image outside the 180° FOV to black.

Regarding the example equipment, the Sigma EX DG 8mm f/3.5 lens was said to have an equisolid-angle projection type by the manufacturer. The real projection function of the lens was measured manually and defined in equation 7.

$$r = 1.97 * f * \sin \frac{\theta}{1.94} \quad (7)$$

where r is the distance on the sensor from the optical axis [mm];

f is the focal length of the lens [mm];

θ is the entrance angle measured from the optical axis [rad].

Since this real projection function is close to the theoretical equisolid projection function, the default distortion function, namely the equisolid-to-equidistant one, is stored in the `fisheye_corr.cal` file.

2.3.5. Determination of the vignetting curves

The vignetting effect in photography is the brightness decrease that can be observed from the centre of a picture toward its periphery, especially when a fisheye lens is used (Reinhard et al. 2006). This vignetting effect depends on the aperture of the lens, with small apertures yielding little vignetting (Inanici 2009). For large apertures, though, the effect can be as large as a 70% luminance loss at the periphery of the fisheye image (Cauwerts et al. 2013). Cauwerts et al.'s and Jakubiec et al.'s studies (Cauwerts et al. 2013; Jakubiec et al. 2016b) suggest that the vignetting curves determined for two devices of a similar brand and model are interchangeable. It is however recommended to derive proper vignetting curves for the device in use since optic materials are not exactly identical, and can contain defects and aberrations (Stanley 2016), and production processes can also change.

The vignetting curve derivation requires evaluating the progressive luminance loss toward the periphery of the fisheye view of an HDR image. For each aperture that will be used later on, a vignetting curve should be determined. There exist at least two common methods of deriving these vignetting curves, which require different equipments.

Vignetting curves derivation from a semi-circular platform. The semicircular platform is made of

evenly distributed uniform grey targets in a room with non-flickering, constant, and indirect lighting (Cauwerts et al. 2013). For each aperture needed, a sequence of multiple exposures (see 2.4.1) of the scene with the targets located at known angles in the FOV of the lens (e.g. every 5°) is captured. The luminance value of the centre of each target is also measured with a spot luminance meter positioned in such a way that its rotating centre is at the same place as the hypothesized no-parallax point of the camera (Fig. 9).

If the lighting stability over time is imperfect, a horizontal illuminance meter is added on the platform to record an illuminance value each time a sequence of multiple exposures is captured. This value is compared to the illuminance value recorded when the spot luminance measures are taken. The spot luminance measures can then be weighted with the percentage difference in horizontal illuminance to account for unstable lighting conditions if need be.

Exposures are selected (see 2.4.2), and HDR images are generated (see 2.4.3) and calibrated in *Photosphere* with the spot luminance measurement made on the central target (see 2.5.6). The vignetting factor for a specific angle of the FOV is evaluated by the ratio of the luminance value extracted from the HDR image to the luminance value measured in the scene for the target corresponding to that angle. Based on the vignetting factors of several angles in the FOV, the vignetting curve, which expresses the vignetting factor that should be applied to a pixel according to its radial position (normalized in a [-1 1] range) in the HDR fisheye image, can be approximated through an even-order polynomial to ensure symmetry (Jacobs 2012). This vignetting curve should be approximated while taking into account to projection type of the HDR image to correct. In this step-by-step process, the projection adjustment is applied before the vignetting correction, hence the projection of the HDR image is equidistant and the radial positions have to be calculated accordingly.

Examples of vignetting factors in function of the normalized radial positions are given for several lens apertures in Table 2. The larger the aperture, the larger the vignetting effect, and the lower the degree of its polynomial.

Vignetting curves derivation by rotation. The second method to derive the lens vignetting




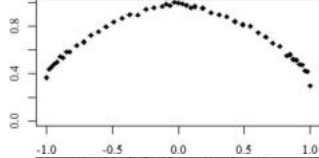
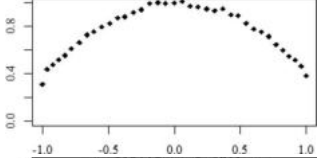
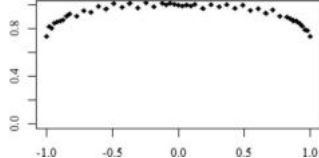
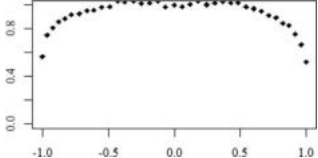
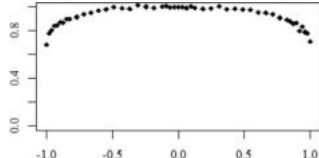
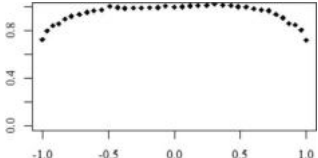
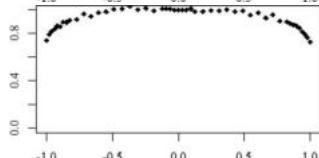
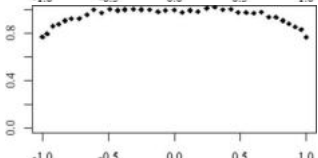
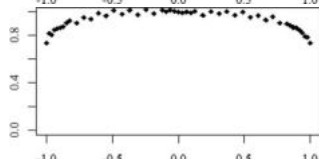
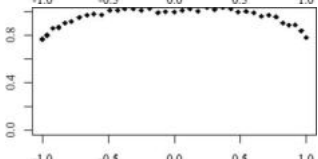
Fig 9 Setup of the equipment for the first method to evaluate the vignetting curves.

curves is more commonly used (Cai and Chung 2011; Hirning 2014; Inanici 2006; Inanici and Galvin 2004; Jakubiec et al. 2016b; Van Den Wymelenberg 2012). It requires a panoramic rotation unit and a constant, non-flickering, and uniform target, which could be either a stable and bright light source in a dark room, or a grey card in a stable lighting environment. The camera is positioned facing the target on a tripod, with the hypothesized no-parallax

point aligned with the centre of rotation of the panoramic rotation unit (Fig. 10).

A sequence of multiple exposures is captured (see 2.4.1). The camera is then rotated by steps (e.g. every 5°) until the camera FOV is covered, and a sequence of multiple exposures is captured at each step. After selecting the exposures (see 2.4.2) and generating the HDR images (see 2.4.3) for each rotation step, the luminance value of the target for each step, and thus each angle of the

Table 2 Comparison of the vignetting curves of the example equipment for different apertures derived with the two suggested methods; in the graphs, the x axis represents the normalized position of the pixel from the centre of the fisheye view [-1 - 1], and the y axis represents the vignetting factor [0 - 1].

f/	Vignetting curves derivation from a semi-circular platform	Vignetting curves derivation by rotation
3.5		
5.6		
11		
16		
22		

FOV, can be extracted. The vignetting factor for a specific angle of the FOV is evaluated by the ratio of the luminance value of the target at that angle (thus derived from the HDR image of the corresponding rotation step) to the luminance value of the target at the centre of the fisheye view image (thus derived from the HDR image facing the target). Similarly to the first method, the vignetting curve is approximated through an even-order polynomial based on the vignetting factors of several angles in the FOV, while accounting for the projection type of the HDR image to correct. This procedure is conducted for each aperture needed.

To counteract the vignetting effect shown by the vignetting curves, the pixels at the periphery of the fisheye view in the HDR image have to be compensated. This is done by multiplying the RGB values of each pixel according to the normalized distance of the pixel from the fisheye view centre, i.e. its normalized radial position, defined in Equation 8, by a factor derived from the inverse of the vignetting curve defined in Equation 9, as shown in Equation 10.

$$z_i = \frac{\sqrt{(x - x_c)^2 + (y - y_c)^2}}{R} \quad (8)$$

$$f_v = \frac{1}{a * z_i^k + b * z_i^{k-1} + \dots + n * z_i + 1} \quad (9)$$



Fig 10 Setup of the equipment for the second method to evaluate the vignetting curves.

$$\begin{aligned} r_o &= f_v * r_i(1); g_o = f_v * g_i(1); \\ b_o &= f_v * b_i(1) \end{aligned} \quad (10)$$

where x and y are the coordinates of the pixel [pixels];

x_c and y_c are the coordinates of the centre of the fisheye view [pixels];

R is the radius of the fisheye view [pixels];

z_i is the distance of the pixel from the optical axis [pixels];

a to n are the parameters of the vignetting polynomial to be defined;

f_v is the vignetting correction factor;

$r_i(1)$, $g_i(1)$ and $b_i(1)$ are the spectrally weighted radiance of the pixel before the vignetting correction [W/m^2sr];

r_o , g_o , b_o are the spectrally weighted radiance of the pixel after the vignetting correction [W/m^2sr].

A .cal text file to apply the vignetting correction is created for each aperture. The .cal files are generated by modifying the extension of a text file (.txt to .cal). Therefore, for each aperture, Equations 8, 9, and 10 are written in a text file, with the parameters of the vignetting polynomial defined according to the aperture, and the extension of the file is then modified to obtain a .cal file readable by *Radiance*. An example of such .cal files is given on p.198 of (Cauwerts 2013).

The vignetting curves for the example equipment were derived using both methods. As seen in Table 2, the two methods provided very similar results, with a maximum of around 65% luminance loss at the periphery of the fisheye image for an aperture of $f/3.5$.

2.3.6. Determination of the ND filter correction function

Pixel overflow, also referred to as sensor blooming, consists of a pixel reaching its limited charge capacity, thus propagating the extra charge to surrounding pixels. To prevent this issue, ND filters should be used when a bright light source having a luminance larger than the upper limit of the equipment luminous range (see 2.3.2) is in the captured FOV. These filters reduce the amount of light hitting the sensor, permitting it to capture larger luminance values. ND filters exist with different density. The transmittance information is given by the ND number according to




Fig 11 Linear regressions estimating the RGB chromatic shifts for the visible portion of the spectrum when the ND filter of the example equipment is used.

Equation 11. For instance, a filter with ND 3.0 has a transmittance of 0.1%, and a filter with ND 0.3, a transmittance of 50%. In lighting research, we recommend using filters with ND 3.0 (daylighting with the sun in the FOV) or lower ND numbers (daylighting and electric lighting).

$$\text{Transmittance}_{\text{ND filter}} = 10^{-\text{ND number}} \quad (11)$$

Although these filters are said to be neutral, they tend to introduce a chromatic shift in the recorded pixel values, especially in the blue channel (Cauwerts 2013). This chromatic shift slightly affects the luminance values of the HDR image (Jakubiec et al. 2016a), since the luminance values are evaluated from the RGB values through a luminance conversion function, as shown in Equations 1 and 2.

Although the correction of the chromatic shift for the entire ND filter spectrum would require the use of a spectroradiometer, the chromatic shift can be corrected for the visible portion of the ND filter spectrum relatively easily. A color transform function which maps the RGB components of an image with the filter to the RGB components of an image without it can be derived and applied to the HDR image to counterbalance the shift (Cauwerts 2013; Inanici 2010; Jones and Reinhart 2017; Stumpf

et al. 2004). For this purpose, a sequence of multiple exposures (see 2.4.1), in which a Macbeth color chart is photographed under stable daylight conditions, has to be captured once with and once without the ND filter. To maximize the radiance range that is used to derive the ND filter correction function, several daylight scenes (e.g. inside and outside) containing the Macbeth color chart should be captured twice (i.e. with and without the ND filter), with the chart in the centre of the fisheye view to avoid the vignetting effect. The exposures are then selected (see 2.4.2) and merged into HDR images (see 2.4.3).

RGB components of each pixel of an HDR image can be extracted using *pvalue*. The line of code to export the RGB components of each pixel in an HDR image (*HDRfile.hdr*) to a text file (*RGB.txt*) using *pvalue* is as follows:

```
> pvalue HDRfile.hdr > RGB.txt
```

Each line of the output text file contains the x and y pixel coordinates followed by the RGB components of the corresponding pixel. The RGB components of each color patch of the Macbeth chart in the HDR images can therefore be derived from this output file.

By comparing the RGB components of each color patch of the Macbeth chart in the HDR images generated with and without the ND filter, the chromatic shift can be estimated. The ND filter correction function is approximated by a linear regression without intercept for each channel (Fig. 11). Besides the chromatic shift correction, this ND filter correction function also takes into account the brightness scaling arising from the use of an ND filter and correct for it. The three linear ND filter correction functions for the three channels are stored in a .cal text file. An example of such a .cal file is given on p.197 of (Cauwerts 2013).

A Kodak Wratten 3.0 ND filter was used with the example equipment when the sun was in the FOV. To derive the ND filter correction function corresponding to each channel, three scenes, including a Macbeth chart, were captured with and without the filter: one inside under an overcast sky, one inside under a clear sky, and one outside under an overcast sky. The three resulting linear functions are shown in Fig. 11.

2.4. HDR image generation

2.4.1. Capture of multiple exposures

To generate an HDR image, a sequence of multiple exposure LDR images has to be captured, with the camera having its settings set as defined in Table 1. This sequence extends from an overexposed image, almost white, to an underexposed image, almost black, and should contain 15 or more LDR images separated by 1 or more EV stop (Fig. 12). The order in which the LDR images are captured could be reversed, and the number of LDR images could be reduced depending on the lighting conditions of the scene.

The exposure variation (1 or more EV stop) between the LDR images is done by changing the exposure time between photographs. Changing the aperture would increase problems associated with vignetting (Reinhard et al. 2006) and shutter speed is a more reliable measure than aperture size (Inanici 2006). For daylighting conditions, it is recommended to avoid exposure times longer than 2 seconds to keep the sequence as short as possible and ensure stable conditions throughout the sequence. The aperture should be set according to the scene, although it is recommended not to use extreme aperture sizes. On the one hand, small aperture sizes (high numerical apertures) are correlated with greater potential for lens flare (Jakubiec et al. 2016b). On the other hand, large aperture sizes (small numerical apertures) suffer from a low maximum captured luminance value (Jakubiec et al. 2016b) and a large vignetting effect (Cauwerts et al. 2013). A mid-range aperture size, such as $f/8$ or $f/11$ for fisheye lenses, can be used as a trade-off.

The camera should be set on a tripod to avoid misalignment problems during the HDR generation (Reinhard et al. 2006). Depending on the number of LDR images to capture, the images can be either saved in jpeg format, which is a compressed file, or in raw format, which is uncompressed and requires around four times

more storing memory. The benefit of raw images is that they contain unprocessed electrical charge information directly from the image sensor (Stanley 2016). Therefore, the generation of HDR images from raw files does not require the preliminary derivation of the response function. *raw2hdr* program (Ward 2011), from the *Radiance* suite, is the equivalent of *hdrgen* for raw images input; all following calibration steps are similar. Although jpeg compression causes the loss of much color data (4096 tonal levels in raw files compared to 256 levels of information for each channel in jpeg files (Stanley 2016)), it is decent at preserving luminance data, and is much more convenient to store. Moreover, when shooting in raw format, the maximum burst during continuous shooting, namely the number of continuous shots that can be taken by the camera without stopping, might be limited to a number lower than the 15 recommended LDR images.

To capture a sequence of multiple exposures, there are three main options. The first one is the manual bracketing, which means that the change of exposure between each image of a sequence is done manually on the camera by modifying the exposure time. This option is not recommended, since the manual adjustment significantly slows down the sequence capture and will most probably cause alignment problems between the images of a sequence. The second option is the automatic bracketing by the camera. This option uses the Automatic Exposure Bracketing (AEB) feature of the camera to directly capture a sequence of multiple exposures. However, not all cameras offer this feature, and most of the time, the maximum number of images that can be captured are limited to 3 or 5. Although it is recommended to capture at least 7 images, 5 images might be sufficient for specific scenes with few contrast. The last option, which is the one recommended, is the automatic bracketing with a remote control. The bracketing settings are defined through a bracketing software,



Fig 12 Sequence of 16 multiple exposure LDR images separated by 1 EV stop captured with the example equipment.

such as *qDslrDashboard* (Hubai 2014), installed on a computer that remotely controls the camera. The capture of multiple exposures is done without interruption and with no risk of movement between the images, and the number of images in a sequence is not limited. Furthermore, in cases where the sun is in the FOV, the use of an ND filter is recommended (see 2.5.6). This filter has to be positioned between the lens and the camera sensor, but it should be handled very carefully. The ND filter has indeed to be kept perfectly clean in order that it adds no light scattering to what is already produced by the lens.

At last, the sequence capture should always be accompanied by at least one spot luminance measurement (Van Den Wymelenberg 2012) and one vertical illuminance measurement (Pierson et al. 2017b). The spot luminance should be measured on a mid-range grey card located in the scene during the sequence capture with a spot luminance meter. This measure will be used at the end of the calibration process to scale the relative luminance values embodied in the HDR image to the absolute luminance values of the scene (see 2.5.6). Furthermore, the vertical illuminance value should also be measured during the sequence capture, and as close as possible to the camera lens but without obstructing it. This illuminance value is necessary to check the validity of the corresponding calibrated HDR image (see 2.5.8).

With the example equipment, times of exposure expanding from 1/8000 to 2 seconds and an aperture of $f/11$ are usually chosen to capture a sequence of 15 multiple exposures separated by 1 EV. The camera, set on a tripod, is remotely controlled through *qDslrDashboard*. Two vertical illuminance measurements are taken, one just before and one just after the sequence, and two or three spot luminance measurements are taken during the sequence on mid-range grey targets in the scene.

2.4.2. Selection of exposures

To ensure that a sequence of multiple exposures contains enough over- and underexposed LDR images and that the entire luminance range of the scene is recorded, it is recommended to capture the largest possible bracketed sequence (Van Den Wymelenberg 2012). However, only LDR images

bringing useful information should be inputted in the merging algorithm in order to accelerate the HDR generation process and make it more stable. Therefore, a selection of LDR images has to be made on the captured sequence. *hdrgen* and *Photosphere* offer a `-x` option to automatically ignore unnecessary exposures. However, it is recommended not to use this option and to make the selection manually since this `-x` option was developed as a timesaver feature and lacks accuracy.

The LDR images have to be selected in such a way that the darkest exposure of the useful sequence has no R, G, and B pixel values greater than 228, and the lightest exposure has no R, G, and B pixel values below 27 (HDRI mailing list as of February 2012). These RGB threshold values are preferred over the 20, 200 RGB thresholds as they have been empirically recognized to achieve better results. The selection rules consist therefore of restraining the sequence from the darkest overexposed image having no pixel value below 27 to the lightest underexposed image having no pixel value over 228. When working on discomfort glare, the second rule is especially important. The darkest exposure having no white pixel means that the total luminance of the brightest light source in the FOV has been captured by the sequence.

More practically, this manual selection could be automated with software such as R (Ihaka and Gentleman 1996). The pixels of an LDR image that are located inside the fisheye view and have a value below 27 (R, G, and B values simultaneously below 27) and over 228 (R, G, and B values simultaneously over 228) are counted. In this way, the LDR images having no white pixel (>228), and those having no black pixel (<27) can be directly identified. In some cases, every overexposed image of a sequence might contain black pixels, or inversely for white pixels. The useful sequence will thus include all over- or underexposed LDR images.

Regarding the sequences of 15 LDR images collected with the example equipment, the selection was automated through an R script on Linux, which consisted of applying a black mask around the fisheye view in the image and counting the pixels before returning the list of selected exposures bringing useful information for the HDR image generation.

2.4.3. Merging of exposures

The merging process is based on the principle that each image in a sequence of multiple exposures will have different pixels properly exposed, but each pixel will be properly exposed in one or more images. Therefore a weighting system is used on the pixels of each image to cut out the influence of over- or underexposed pixels and increase the influence of correctly exposed ones in the merging of the exposures.

In summary, the HDR generation process from jpeg files consists first of linearizing the pixel values in each LDR image by canceling out the camera response function. Then, the pixel values of each LDR image are divided by the image exposure time to bring the pixel values into the same domain. At last, each pixel value of each LDR image is weighted according to its value, and the resulting pixel values of the HDR image are the averages across exposures of the weighted corresponding pixel values (Reinhard et al. 2006).

hdrgen or *pfs calibration* are examples of merging algorithms that require the camera response curve and a sequence of multiple exposures to generate an HDR image. The choice between these two tools should have no impact on the quality of generated HDR images since both Debevec and Malik's and Robertson's algorithms have been proved to perform equally well (Jacobs and Wilson 2007). However, other HDR merging software such as *Bracket*, *Luminance HDR*, and *Picturenaut*, have been shown to produce poorer results (Stanley 2016) and should therefore be avoided. Moreover, if a proper camera response function has not been previously derived as recommended (see 2.3.3), one will be created by the algorithm based on the inputted sequence of multiple exposures. Not deriving a proper camera response function will most probably cause a larger error in the HDR image since the captured scene that is then used to derive a substitute response function was not optimized.

hdrgen offers some features amongst which the main ones are the automatic exposure alignment (-a), the exposure adjustment (-e), the lens flare removal (-f), the ghost removal (-g), and the under- or overexposed image removal (-x). When taking a tripod-stabilized sequence, the automatic

exposure alignment (-a) feature should be turned off to avoid unnecessary manipulations of the HDR image. Similarly, for auto-bracketed sequence, the exposure adjustment (-e) feature should be avoided, especially if the camera reproducibility test was successful (see 2.1.2). This feature adjusts for errors in the camera exposure settings by fine-tuning the scalar multiplier between exposures. Moreover, if the captured scene is prone to lens flare, namely scattered light from bright light source which results in a slightly fogged or starred appearance in the image (Fig. 13), or ghosts, i.e. objects or subjects moving during the sequence, the corresponding feature (-f and -g) should be activated in *hdrgen* to remove these artifacts.

In any case, the selection of LDR images should not be executed through the image removal (-x) feature to avoid removing images that could bring useful information. The line of code to merge the selected LDR images (*LDRfiles_*.jpg*) into an HDR image (*HDRfile.hdr*) with the predefined camera response function (*response_function.rsp*) using *hdrgen* is as follows:

```
> hdrgen LDRfiles_*.JPG -o HDRfile.hdr -r response_function.rsp -a -e -f -g
```

HDR images generated by *hdrgen* or *Photosphere* include a header, where the settings and parameters of the image are stored (Jacobs 2007). The line of code to visualize the header information of an HDR image (*HDRfile.hdr*) using *getinfo* is as follows:

```
> getinfo HDRfile.hdr
```

The header information includes the view type, the viewing angle and the exposure value. It must be noted that the view type and viewing angle written in the header automatically are typically not correct, since this information cannot be retrieved from the EXIF information of the LDR images. Therefore these header entries have to be modified to the correct values later on (see 2.5.7). In addition, some programs of *Radiance* will interact with or deactivate the header entries when applied to the HDR image. This is the case for instance with *pcomb* and *pcompos*, which will mark the exposure or view type entries in the header as "invalid" by adding a tabulator. The




Fig 13 Example of visible lens flare due to the sun in the field of view.

data contained in an HDR image will therefore be wrongly interpreted since by default, an exposure value of 1 is used when the exposure value cannot be extracted from the header. It is therefore required to nullify the exposure entry in the header (see 2.5.1).

2.5. HDR image calibration

2.5.1. Nullification of exposure value

To prevent the issue of wrong interpretation of the exposure, it is recommended to include the exposure value directly in the pixel values, and remove the exposure entry completely. The recommended method is to use *ra_xyze*, which includes the exposure in the pixels and removes it from the header, as the first command applied after an HDR image

is generated. The line of code to nullify the exposure value of an HDR image (*HDRfile.hdr*) using *ra_xyze* is as follows:

```
> ra_xyze -r -o HDRfile.hdr > HDRfile_EVinpixel.hdr
```

2.5.2. Cropping and resizing

To facilitate the manipulations applied to an HDR image, it is recommended to crop this image to a square encompassing the circular fisheye view (Fig. 14).

Using the fisheye view coordinates defined during the one-time setup (see 2.3.1), the cropping of the HDR image is carried out using *pcompos*. The input values required are the coordinates (in pixels) of the bottom-left corner of the square encompassing the total viewing angle, as well as the fisheye view diameter (in pixels). The line of code to crop the HDR image (*HDRfile.hdr*) to a smaller pixel resolution (*diameter*) having the anchor point set as the bottom-left corner of the square encompassing the fisheye view (*xleft*; *ydown*) using *pcompos* is as follows:

```
> pcompos -x diameter -y diameter HDRfile.hdr -xleft -ydown > HDRfile_cropped.hdr
```

By applying this command, the header entries for the view type and the viewing angle are marked as “invalid” and need to be adjusted later (see 2.5.7).

Moreover, there might be some applications for which the size of an HDR image is limited to a certain pixel resolution in order to avoid unreasonable computational processing time. This is the case for discomfort glare analyses with *evalglare*, since the computational performance drops when processing HDR images larger than 1500×1500 pixels. The resizing of




Fig 14 Cropping of an HDR image to a square encompassing the circular fisheye view.

an HDR image, which is optional and dependent on the final use of the HDR image, can be done through *pfilt* with the HDR image and the desired size in pixels as inputs. While using *pfilt*, the *-1* feature should be activated to allow only one pass on the file and avoid the addition of an exposure entry in the header of the HDR image. The line of code to resize the HDR image (*HDRfile.hdr*) to a smaller resolution (*xdim*, *ydim*) using *pfilt* is as follows:

```
> pfilt -1 -x xdim -y ydim HDRfile.hdr >
HDRfile_resized.hdr
```

Cropping and/or resizing an HDR image implies that the new size and new centre coordinates have to be considered in the *.cal* files that have been determined during the one-time setup (e.g. the *.cal* files for the vignetting correction).

2.5.3. Projection adjustment

In applications for which the geometry of the visual scene is relevant, it is often preferred to have an HDR image with one of the two most common theoretical projections (i.e. equidistant or orthographic). If the real projection function of the lens differs from these projection types, the projection of the HDR image can be adjusted thanks to the distortion function stored in the *fish-eye_corr.cal* file (see 2.3.4). The line of code to apply the projection adjustment to an HDR image (*HDRfile_equisolidproj.hdr*) using *pcomb* and the *fish-eye_corr.cal* file (*fish-eye_corr.cal*) containing the correct distortion function (by default, the distortion function encoded is the equisolid-to-equidistant distortion) is as follows:

```
> pcomb -f fish-eye_corr.cal HDRfile_equisolidproj.
hdr > HDRfile_equidistantproj.hdr
```

Besides adjusting the projection of an HDR image, the *fish-eye_corr.cal* file also sets the pixel values outside the 180° fisheye view to 0 for all three channels, namely to black.

2.5.4. Vignetting correction

To correct the light fall-off toward the edges of an image due to the vignetting effect, the *.cal* text file containing the vignetting correction function corresponding to the lens aperture is used (see 2.3.5). The vignetting correction can be applied to an

HDR image by using *pcomb* with the correct *.cal* file. The line of code to apply the vignetting correction (*vignetting_fX.cal*) corresponding to the aperture used for the input HDR image (*HDRfile.hdr*) using *pcomb* is as follows:

```
> pcomb -f vignetting_fX.cal HDRfile.hdr >
HDRfile_vignettingcorrected.hdr
```

If the pixels outside the fisheye view have not been previously set to black (R, G, B = 0, 0, 0), the vignetting correction function might apply to those as well, and alter the aspect of the HDR image. This has no impact on the data contained in the fisheye view, and the pixels outside of the fisheye view can be reset to black afterward. This black setting is not needed in case of use of *evalglare*.

2.5.5. ND filter correction

When an ND filter is used to capture a sequence of multiple exposures, the resulting HDR image needs to be corrected for the chromatic shift and the brightness scaling introduced by the filter (see 2.3.6). The correction of the ND filtered HDR image can be done with *pcomb*, which applies the three linear correction functions (one for each RGB channel) stored in the *.cal* text file to the pixel values. The line of code to apply the ND correction (*ND_correction.cal*) to a filtered HDR image (*HDRfile.hdr*) using *pcomb* is as follows:

```
> pcomb -f ND_correction.cal HDRfile.hdr >
HDRfile_NDfiltercorrected.hdr
```

2.5.6. Photometric adjustment by spot luminance measure

HDR photography allows capturing relative luminance values of all the pixels in the lens FOV. In order to retrieve the absolute luminance values of the scene, a photometric adjustment has to be applied to the HDR image, which consists of multiplying the relative luminance values of every pixel by a scaling factor. This scaling factor is the ratio of the measured to the HDR-derived luminance value of a target in the FOV (Jacobs 2007). The factor is believed to be dependent on the lens aperture and the type of the main lighting source, and should

theoretically oscillate around 1 (Brotas and Jacobs 2006). In practice, there exist two main methods to adjust an HDR image photometry with spot luminance measures, depending on the conditions under which the HDR images are captured.

When HDR images of different scenes are recorded, it is recommended to apply a proper scaling factor to each generated HDR image (Van Den Wymelenberg 2012). The photometric adjustment therefore requires the measurement with a spot luminance meter of the luminance value of a target in the FOV, such as a mid-range grey card (Fig. 15), during each capture of multiple exposures.

The grey color of the target should be as neutral as possible to reflect evenly all wavelengths between 380 nm and 780 nm, and should have a mid-range reflectance to be representative of the visual scene. Mid-range grey targets have been shown to produce lower errors in terms of luminance than colored, black, or white targets (Inanici 2006). Furthermore, the spot luminance measure should be done in an area of the scene with mid-range luminance values, and especially not in an area prone to lens flare or pixel overflow. Since the exact pixel coordinates of the target in the fisheye view are unknown a priori, the photometry adjustment has to be done manually, for instance through *Photosphere*. To do so, the target has first to be found on the displayed HDR image in *Photosphere* viewer. Then the calibration is automatically applied to every pixel of the HDR image by inputting the measured spot luminance value for the target area.

When several HDR images are generated for the same scene under similar lighting conditions, a calibration factor can be derived for the first HDR image and applied to all subsequent images generated with the same equipment and the same lens aperture. There is therefore no need to collect a spot luminance value on a mid-range grey target in the FOV for each HDR image, but only for each first HDR image using a specific aperture. The calibration factor corresponding to that specific aperture is then defined as the ratio of the measured to the HDR-derived luminance value of the target. The photometric adjustment is applied to the HDR image using *pcomb* with the calibration factor corresponding to the



Fig 15 Example of a visual scene with two mid-range grey cards used as targets for spot luminance measurement.

HDR image aperture. The line of code to apply the calibration factor (K) to an HDR image (*HDRfile.hdr*) using *pcomb* is as follows:

```
> pcomb -s K HDRfile.hdr > HDRfile_CF.hdr
```

For the example equipment, the mean calibration factors for indoor scenes with daylight as the principal lighting source were defined as 1.21 for an aperture of f/8, and 1.24 for an aperture of f/11. For indoor scene under electric lighting, the mean calibration factor could reach 1.87 for an aperture of f/22.

2.5.7. HDR image header editing

Correct information in the HDR image header is essential for quite all post-analyses. Besides the calculation of metrics, which are generally depending on correct angular relations in the image (e.g. glare metrics), the calculation of an average value of a zone requires the consideration of the solid angle of each pixel to be mathematically exact. Since the view information stored in the header is not correct (see 2.4.3), postprocessing tools will calculate the angles (including the solid angles) incorrectly, distorting the results. For that reason, the view information has to be adjusted.

The view information in the header consists of the projection type of the lens, as well as the total viewing angle of the lens. The projection type is reported in the header as *-vta* for an equidistant projection, *-vth* for an orthographic projection, or *-vtv* for a non-fisheye lens. The total viewing angle is reported in degrees both vertically, i.e. *-vv*, and horizontally, i.e. *-vh*.

The first step to adjust the view information in the HDR image header is to erase the existing information. The line of code to erase the view entry from the header of an HDR image (*HDRfile.hdr*) using *getinfo* is as follows (under Windows, the line of code requires the GNU package “*sed for windows*” to be installed):

```
> (getinfo < HDRfile.hdr | sed "/VIEW/d" &&
getinfo - < HDRfile.hdr) > HDRfile_withoutView.hdr
```

In a second step, the correct view information has to be inserted back in the HDR image header. The line of code to set the correct view information (*VIEW = -vt? -vv X -vh X*”) in the header of an HDR image (*HDRfile.hdr*) using *getinfo* is as follows:

```
> getinfo -a VIEW= -vt? -vv X -vh X" <
HDRfile.hdr > HDRfile_correctView.hdr
```

Furthermore, it was recommended in a previous step (see 2.5.1) to include the exposure value in the pixels and remove it from the HDR image header. By doing so, the header contains no exposure information, implying a default exposure value of 1.

Since an HDR image generated with the example equipment and calibrated according to the previous steps has an equidistant projection and a total viewing angle of 186°, the view information in the HDR image header is changed to “VIEW= -vta -vv 186 -vh 186”.

2.5.8. Validity check

At the end of the process, each HDR image should be checked before being used as a luminance map. The reason for this is that even following strictly the described step-by-step process, there might occur unexpected significant errors in the HDR images. Therefore, especially for use in research, each HDR image has to undergo a quality check before being used. This check should also be reported when results based on HDR images are published.

There exist at least two ways for checking the validity of the luminance measurements made with HDR photography: the illuminance comparison and the luminance comparison. The illuminance check should be preferred, especially when expecting high luminance (>30,000 cd/m²) in the HDR image (e.g. sky, sun or reflection of the sun, sun visible through shadings, etc.). While the luminance check only looks at a restricted reference area, the illuminance check considers the integral of the luminance values over the whole hemisphere and is therefore more thorough, and more appropriate to detect pixel overflow.

Illuminance comparison. The illuminance comparison checks the luminous balance to ensure that the total lighting energy embodied in the HDR image is equivalent to the total lighting energy reaching an illuminance sensor placed just above the lens. More specifically, the integration of all luminance values over the hemispheric FOV captured in the HDR image should correspond to the illuminance value measured by a properly calibrated illuminance meter placed next to the lens (Kent 2016). By positioning the illuminance meter above the fisheye lens, it is assured that, in the unfortunate situation where the lens is shading part of the illuminance meter sensor, it only shades the light reflected from the ground. On the opposite, if the illuminance meter was shading part of the fisheye view, this would be directly visible in the image.

In practice, the integration of all luminance values over the hemispheric FOV captured in the HDR image corresponds to the illuminance of the calibrated HDR image. The illuminance of the HDR image can be derived through *evalglare* and its *-V* option (Wienold et al. 2004), and compared to the illuminance reading of the sensor. The line of code to retrieve the illuminance value of an HDR image (*HDRfile.hdr*) using *evalglare* is as follows:

```
> evalglare -V HDRfile.hdr
```

When the error between the sensor-measured and HDR-derived illuminance values is larger than 25% (see 3.1) although the calibration process was conducted meticulously, it most probably means

that for some areas of the HDR image, the luminance is over- or under-estimated. In the specific case where the measured illuminance value is significantly larger than the HDR-derived value and there is a very bright spot in the FOV such as the sun, the luminance underestimation might be explained by pixel overflow in the HDR image (see 2.3.6). In case pixel overflow happened in recorded images and a recapturing with an ND filter is not possible, *evalglare* offers a correction function (Wienold and Andersen 2016), and alternatives have also been proposed by Jakubiec et al (Jakubiec et al. 2016a).

Luminance comparison. An additional way of checking the validity of calibrated HDR images is to take more than one spot luminance measure in the scene when capturing the sequence of multiple exposures, especially on low- or high-range grey card placed in the dark and bright parts of the scene. These measured values can then be compared to the HDR-derived values retrievable with *ximage* or *Photosphere* for instance. The line of code to display an HDR image (*HDRfile.hdr*) using *ximage* is as follows:

```
> ximage HDRfile.hdr
```

Once the HDR image is displayed via *ximage*, the luminance can be retrieved by clicking on the point of interest in the HDR image viewer and typing *l* (like luminance).

For scenes in which the luminance range does not exceed $30,000 \text{ cd/m}^2$, one could expect less than 10% error, as it has been demonstrated in the literature (Cai and Chung 2011; Inanici 2006). However, the luminance error from HDR images of more extreme scenes cannot be extrapolated from these results, and further research is needed to define a quality threshold for luminance error of HDR images with bright light sources.

In brief, Table 3 summarizes the steps necessary to generate and calibrate a 180° HDR image of a daylight scene from a sequence of multiple exposures in jpeg format, as well as the material required to execute these steps, which is either measured or has been derived during the one-time setup, and the main *Radiance* commands. For the command of step 10, *vta* needs to be changed to *vth* for an hemispheric projection.

3 Discussion

3.1. Expected range of errors

A dataset of 795 HDR images was used to evaluate the range of errors that can be expected from HDR photography and this step-by-step calibration process. This dataset was gathered during a field study conducted in four countries (Chile, Belgium, Japan, and Switzerland) from March 2017 to August 2018. 180° luminance maps containing very diverse office scenes, namely workers' FOV when sitting at their desk, were collected for the study. The study took place under clear sky conditions in different office buildings for which the main light source was daylight, and for office desks that were located next to an east-, west-, or south-oriented (north- for Chile) window. The detailed experiment protocol is available in (Pierson et al. 2017a).

The LDR images were all captured with the example equipment. The previously described step-by-step process was applied rigorously to generate and calibrate, with a proper calibration factor for each image, the 795 luminance maps. From this dataset, the mean error in terms of vertical illuminance, with the error being calculated according to Equation 12, for all luminance maps is 3.05%, and the largest error is 74.88%. Only 1.38% of the luminance maps (11 maps) showed an error of more than 25%.

$$E_v[\%] = \frac{|E_{v,HDR} - E_{v,measured}|}{E_{v,measured}} * 100 \quad (12)$$

where E_v is the error in vertical illuminance [%];

$E_{v,HDR}$ is the HDR-derived vertical illuminance [lux];

$E_{v,measured}$ is the sensor-measured vertical illuminance [lux].

Because it is expected that luminance maps having a large range of luminance will be prone to larger errors, a detailed error analysis (Table 4) is performed separately for the luminance maps of the dataset having a maximal luminance value lower than $30,000 \text{ cd/m}^2$, and those having a maximal luminance value larger than $30,000 \text{ cd/m}^2$.

From Table 4, it is observed that the errors in vertical illuminance are larger for luminance maps having a high luminance range compared to

Table 3 Summary of the steps, required material (measured or derived), and *Radiance* commands to generate and calibrate a 180° HDR image of a daylight scene from a sequence of multiple exposures in jpeg; the part of the commands written in italic text need to be adjusted to the real values.

Calibration step	Required material	Generation and calibration command (Linux)
(1) Capture of exposures	/	This step can be automated through different software (qDslrDashboard, digiCamControl, etc.)
(2) Selection of exposures	/	This step can be automated through different software (R, python, etc.)
(3) Merging of exposures	multiple exposures +	<code>hdrgen -JPG -o <i>output1.hdr</i> -r <i>camera.rsp</i> -a -e -f -g</code>
(4) Nullification of exposure value	response function	
(5) Cropping and resizing	/	<code>ra_xyze -r -o <i>output1.hdr</i> > <i>output2.hdr</i></code>
(6) Projection adjustment	distortion function	<code>pcompos -x <i>diameter</i> -y <i>diameter</i> <i>output2.hdr</i> -xleft -ydown > <i>output3.hdr</i></code>
(7) Vignetting correction	vignetting curves	<code>pfilt -1 -x <i>xdim</i> -y <i>ydim</i> <i>output3.hdr</i> > <i>output4.hdr</i></code>
(8) ND filter correction	ND filter correction function	<code>pcomb -f <i>distortion.cal</i> <i>output4.hdr</i> > <i>output5.hdr</i></code>
(9) Photometric adjustment	measured L_T or calibration factor	<code>pcomb -f <i>vignetting.cal</i> <i>output5.hdr</i> > <i>output6.hdr</i></code>
(10) HDR image header editing	projection type +	<code>pcomb -f <i>correction.cal</i> <i>output6.hdr</i> > <i>output7.hdr</i></code>
	real viewing angle	<code>pcomb -s <i>factor</i> <i>output7.hdr</i> > <i>output8.hdr</i></code>
		<code>(getinfo < <i>output8.hdr</i> sed "VIEW/d" && getinfo - < <i>output8.hdr</i>) > <i>output9.hdr</i></code>
		<code>getinfo -a "VIEW= -vta -vv <i>angle</i> -vh <i>angle</i>" < <i>output9.hdr</i> > <i>output10.hdr</i></code>
(11) Validity check	measured E_v +	<code>evalglare -V <i>output10.hdr</i></code>
	luminous range	

luminance maps having a lower luminance range. In general, an error in vertical illuminance of less than 10% should be expected, and HDR images with an error of more than 25% should be rejected. The numbers given in Table 4 are derived from the absolute value of the errors, thus without considering their \pm sign. But it should be noted that a systematic positive (or negative) error in vertical illuminance should not occur. If the mean error (not in absolute terms) is not approaching zero, it means that an error causing a systematic increase (or decrease) in vertical illuminance was inserted in the calibration process.

Furthermore, although it is recommended to apply all the steps of the here-described calibration procedure and a proper calibration factor for each HDR image, it might happen that this procedure cannot be executed in its entirety. It is therefore interesting to know the dimensions of the error that should be expected if a specific step is not performed or if proper calibration factors cannot be derived.

Regarding the photometric adjustment of the HDR images (see 2.5.6), error statistics have been evaluated for the low- and high-range luminance maps of the dataset (Table 5), to compare the error ranges when:

individual calibration factors (CFs) are used for each image;

a predefined CF is used for all images with the same aperture;
no CF is derived, and no photometric adjustment is applied.

As expected, the use of individual CFs provides the most accurate luminance maps in terms of vertical illuminance, while not applying any photometric adjustment causes only 25% of the luminance maps to have an error lower than 10%. However, the errors calculated for the luminance maps without any photometric adjustment are very dependent on the equipment in use, and another equipment might give a different error range.

Since the errors in vertical illuminance related to the other steps of the described process are highly dependent on the scene being evaluated, the error statistics derived from the dataset for each missing step are not reported. For instance, a dataset made of extreme scenes with very bright light sources in an off-centre position most probably suffer from larger errors in vertical illuminance due to the omission of the vignetting correction than a dataset mostly made of quite uniform indoor daylight scenes.

Instead, a description of the impact that could be brought by the omission of each specific step of the process is made in Table 6. An estimation

Table 4 Vertical illuminance errors evaluated from a set of luminance maps collected with the example equipment and calibrated with the described step-by-step calibration process applying a proper calibration factor for each image.

	HDR images with $L_{\max} < 30,000\text{cd/m}^2$	HDR images with $L_{\max} > 30,000\text{cd/m}^2$	All HDR images
N	666	129	795
Mean error	2.5%	6.1%	3.1%
Maximal error	37.7%	74.9%	74.9%
HDR images with error < 25%	663 (99.6%)	121 (93.8%)	784 (98.6%)
HDR images with error < 10%	648 (97.3%)	112 (86.8%)	760 (95.6%)

Table 5 Comparison of vertical illuminance errors between the different photometric adjustment methods for the set of luminance maps collected with the example equipment and calibrated with the described step-by-step calibration process.

	HDR images with $L_{\max} < 30,000\text{cd/m}^2$			HDR images with $L_{\max} > 30,000\text{cd/m}^2$		
	Individual CFs	Predefined CFs	No CFs	Individual CFs	Predefined CFs	No CFs
Mean error	2.5%	6.6%	14.5%	6.1%	12.1%	18.6%
Maximal error	37.7%	55.0%	56.0%	74.9%	131.9%	96.5%
HDR images with error < 25%	99.6%	99.4%	95.0%	93.8%	89.9%	84.5%
HDR images with error < 10%	97.3%	79.9%	23.7%	86.8%	61.2%	25.6%

based on theoretical principles of the error due to the omission of each step is also given.

Furthermore, it should be noted that, up to now, there exist no solutions to correct for lens flare, although several authors have acknowledged or quantified the related error (Anaokar and Moeck 2005; Inanici 2006; Rea and Jeffrey 1990). Lens flare is the radial scattering of light information of a single pixel into neighboring pixels (Fig. 13), which results in a loss of luminance information of the central pixel and an increase of false luminance data in the neighboring pixels (Stanley 2016). Therefore, it should be reminded that the surrounding area of a bright light source in an HDR image is most probably prone to an error in luminance values due to lens flare.

Even when the complete calibration procedure is performed and a proper calibration factor is applied, some of the luminance maps have an error larger than 25% in vertical illuminance, which cannot be explained by pixel overflow. This is the case for 0.5% ($n = 4$) of the luminance maps of this dataset. Up to now, this error cannot be explained and these images have to be disregarded. If an illuminance comparison check had not been conducted for every calibrated HDR images, it might have happened that these erroneous luminance maps corrupt the dataset, making the associated results misleading. This validity check is also a good way of detecting when one or

several steps of the calibration procedure went wrong, causing all the luminance maps to be flawed.

3.2. Regarding electric lighting studies

In the case of electric lighting, the HDR image generation and calibration process might differ in several ways to the process described in this tutorial. The main differences and challenges that should be accounted for when creating luminance maps of electrically lit scenes are highlighted below.

When the predominant lighting type is modified, the first adaptation of the process is the modification of the white balance setting of the camera if LDR images are recorded in jpeg format. The white balance setting should be defined according to the color temperature of the studied lighting conditions. A recent study showed that, for a scene in which the predominant lighting source was fluorescent, a daylight white balance setting produced response curves with extreme abnormalities and inconsistencies whereas a cool-white fluorescent white balance setting produced very consistent response curves (Stanley 2016).

Secondly, the derivation of the response function and the predefined calibration factors, as well as the chromatic correction if an ND filter is used, should all be carried out in similar lighting conditions than the ones that will be measured.

Table 6 Impact and estimated error that can be expected when one of the steps of the calibration process is not executed (DGP = Daylight Glare Probability; L_{avg} = average Luminance).

Missing step	Impact	Consequence
Selection of exposures	Superfluous images used in the merging step leading to more weight given to not-properly exposed pixels	Error in luminance in the range of 0-5% for some pixels
Nullification of exposure value	Loss of exposure information of the HDR image	Error in luminance up to 100% for all pixels
Projection adjustment	Wrong calculation of angles and solid angles during HDR image post processing	Error in solid angle size up to 20% for some pixels (e.g. for equisolid-angle lenses)
Vignetting correction	Underestimation of the luminance in the outer angle ranges of the lens field of view	Depending on lens and aperture, error in luminance up to 70% for the pixels in the outer angle ranges
Photometric adjustment	Non-scaled relative luminance values	Error in luminance up to 100% for all pixels
Header editing: view type	Wrong calculation of angles and solid angles during HDR image post processing	Error in vertical illuminance and in metrics using angles in their equation (i.e. DGP, L_{avg}) up to 100% (e.g. when a non-fisheye view type is used for a fisheye lens)
Header editing: viewing angle	Wrong calculation of angles and solid angles during HDR image post processing	Error in vertical illuminance and in metrics using angles in their equation (i.e. DGP, L_{avg}) in the range of 0-5%

Furthermore, other issues will have to be taken care of depending on the kind of electric light source being measured through the HDR photography.

When light sources operated on alternating currents are the main lighting component of the scene, the flicker, which is a repetitive change in brightness of an electric light source generally invisible to the human eye due to its rapid frequency (IESNA 2000), might be one of these issues. Since multiple exposures are captured, the lighting conditions of the scene have to remain constant throughout the sequence. The flicker might therefore interfere with the recorded values during the short exposures. For instance, it has been shown that the error related to the flicker of a LED could be larger than 40% for an exposure time of 1/250s (Jacobs 2012).

To measure the flicker frequency, an oscilloscope is needed. If the flicker frequency is known or measured, the shortest exposure time should be 10 times the flicker frequency to reduce the impact of flicker to a minimum. But if the flicker frequency cannot be measured, shooting a series of 100 images with the same exposure time and fixed settings, and comparing the differences in pixel values between the images gives an indication of the error related to the flicker that can occur for that exposure time. After repeating this procedure for different exposure times (e.g. 100 images for 1/500s, 100 images for 1/250s, etc.), the shortest exposure time can be defined

according to what is acceptable as an error related to flicker.

To be on the safe side, and without measuring the flicker frequency or doing any additional tests, the following rule of thumb should be used: the shortest exposure time should be 10 times the grid-net frequency (assuming the lowest possible flicker frequency equals the electrical grid frequency). This implies that for a 50 Hz grid, the shortest exposure time should be 0.2 s (1/5 s). In this case, collecting a series of 15 under- and over-exposed images requires the use of an ND filter.

Another potential issue lies with the resolution of HDR images and the size of the light source(s) being measured. For instance, current LED luminaires are designed as a juxtaposition of very small LED light sources, which could induce a contrast to the human eye depending on several factors such as the distance. However, due to its resolution still very limited compared to our own visual system, a camera might not be able to capture this contrast. Therefore, the recorded luminance distribution of the luminaire will not be representative of the human eye perception as the peak luminance of the individual LED light sources will be averaged in the pixels of the luminaire area.

These attention points do not constitute an exhaustive list, and other challenges might emerge when measuring electrically lit scenes through HDR photography with semiprofessional equipment.

3.3. Regarding color and circadian lighting studies

The procedure detailed in this paper aims at ensuring photometric accuracy, but does not necessarily lead to colorimetric accuracy. For an application in circadian lighting or color studies, the generation and calibration of HDR images should be executed through an alternative procedure.

Jung and Inanici (Jung and Inanici 2018) recently investigated HDR photography for measuring circadian lighting. They proposed a procedure in which CIE XYZ values should first be retrieved using the standardized sRGB to CIE XYZ color transformation matrix (IEC 1999). In a second step, each channel should be adjusted with a specific calibration coefficient determined by comparing tristimulus values (XYZ values) measured with an illuminance based color meter with the triplet retrieved from the HDR fisheye image. From these color calibrated images, equivalent melanopic illuminance values could be derived. An average error of 10% and a maximum error of 16% were observed. It was concluded that HDR fisheye images (generated from a Canon EOS 5D fitted with a Sigma 8mm F3.5 EX DG) can be used to measure circadian lighting values (per-pixel equivalent melanopic luminance and illuminance values).

Alternatively, rather than making the assumption that camera sensor primaries are ITU-R BT.709 primaries and using the sRGB transform matrix for retrieving CIE XYZ values, a color transform matrix can be determined specifically for each camera/lens association (Kim and Kautz 2008).

4 Limitations

It should be noted that semiprofessional DSLR-cameras are no measurement tools. HDR photography has however been shown as a reliable technique to measure luminance (Inanici 2006). Therefore, the step-by-step calibration process described here can be used for typical indoor daylight scenes, although larger luminance errors could be expected for darker areas in a visual scene.

5 Conclusion

A synthetic and pragmatic overview of the step-by-step calibration process to generate a calibrated HDR fisheye image with semiprofessional equipment is presented. The described process is aimed at creating measured luminance maps for daylighting studies (thus for daytime applications), although the main attention points to adapt the process for electric and circadian lighting studies are highlighted.

All the presented calibration steps, as well as the one-time setup steps, are necessary to produce an accurate luminance map. Along the process, examples established with a Canon EOS 5D Mark II camera, a Sigma EX DG 8mm f/3.5 fisheye lens, and the *Radiance* suite of programs (Ward 1987) are given.

Besides *Radiance*, more user-friendly computer software products offer to generate and calibrate HDR images from a sequence of multiple exposures. To this date, none of the most common ones carry out all the steps described in this tutorial. For instance, *hdrscope* misses the HDR image generation, the optional ND filter correction and the projection adjustment; *Photosphere* misses the optional ND filter correction, the vignetting correction, and the projection adjustment; *WebHDR* misses the HDR image cropping, the optional ND filter correction, the vignetting correction, and the projection adjustment. Another downside of these software products is the inability of the user to fully understand the modifications applied to the images, and therefore to detect a potential error along the process.

A common pitfall of HDR calibration is the use of outdated versions of *Radiance* or any other software. Freely available software tools are constantly being tested by the community, and corrected when need be. Therefore, it is important to make sure that the version in use is the latest release, and avoid introducing problems which have already been solved in a most recent version.

Moreover, although it would allow many more people to create luminance maps through HDR photography, the use of tablets, smartphones, and any other connected devices to capture a sequence of multiple exposures is not recommended. Studies (Bieckmann 2016; Cerqueira et al. 2018) have indeed

shown a lack of reproducibility with this kind of equipment, especially if updates are being installed on the device without the user knowing about it.

At last, the given examples show that a high measurement accuracy and a good reproducibility can be obtained from HDR photography with commercially available equipment when the calibration process is meticulously applied. For luminance maps with luminance values lower than 30,000 cd/m², 97% of calibrated HDR images generated with the example equipment demonstrate an error in vertical illuminance lower than 10%, against 87% for luminance maps with a higher range of luminance. However, a few calibrated HDR images still present an error in vertical illuminance of up to 75%. Some of these large errors can be explained by pixel overflow, though not all of them. The validity check based on an external measurement, such as the vertical illuminance, should therefore be mandatory for each generated luminance map in lighting research, and should be reported in related publications. Luminance maps failing the quality measure should be removed from the evaluated dataset.

In the future, a calibration process for the HDR photography with commercially available equipment should be established for electric lighting, road lighting, and color studies. This could be the opportunity for the lighting community to standardize the calibration process for HDR photography and define a quality criterion that should be reported and met by all lighting studies using this measurement technique.

Acknowledgments

The authors would like to thank Axel Jacobs and Greg Ward for their comments and suggestions on the content of this article.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

Clotilde Pierson and Coralie Cauwerts are respectively Research Fellow and Postdoctoral Researcher of the Fonds de la Recherche Scientifique – FNRS. This article is published with the support of the Fondation Universitaire de Belgique.

ORCID

C. Pierson  <http://orcid.org/0000-0001-7847-6568>
 C. Cauwerts  <http://orcid.org/0000-0001-9107-0810>
 M. Bodart  <http://orcid.org/0000-0003-2826-6751>
 J. Wienold  <http://orcid.org/0000-0002-3723-0323>

References

- Anaokar S, Moeck M. 2005. Validation of high dynamic range imaging to luminance measurement. *LEUKOS*. 2(2):133–144.
- Anderson H. 2006. The nodal point. [accessed 2019 Jan 27]. <http://www.hugha.co.uk/NodalPoint/Index.htm>.
- Bieckmann K. 2016. Measuring illuminance via smartphone? Paper presented at: LICHT 2016. Karlsruhe, Germany.
- Brotas L, Jacobs A. 2006. Webhdr. 5th International Radiance Workshop. Leicester, UK.
- Cai H, Chung TM. 2011. Improving the quality of high dynamic range images. *Light Res Technol*. 43:87–102.
- Cauwerts C. 2013. Influence of presentation modes on visual perceptions of daylight spaces [PhD dissertation]. Université catholique de Louvain (UCL).
- Cauwerts C, Bodart M, Deneyer A. 2013. Comparison of the vignetting effects of two identical fisheye lenses. *LEUKOS*. 8(3):181–203.
- Cerqueira D, Carvalho F, Melo RB. 2018. Is it smart to use smartphones to measure illuminance for occupational health and safety purposes? Paper presented at: Intl Conference on Applied Human Factors and Ergonomics. Orlando, FL.
- Debevec PE, Malik J. 1997. Recovering high dynamic range radiance maps from photographs. Paper presented at: SIGGRAPH 97. Los Angeles, CA.
- Dumortier D, Coutelier B, Faulcon T, Van Roy F. 2005. Photolux: A new luminance mapping system based on nikon coolpix digital cameras. Paper presented at: Lux Europa 2005. Berlin, Germany.
- Geisler-Moroder D, Lee E, Ward G. 2016. Validation of the radiance 5-phase-method against field measurements. Paper presented at: 15th International Radiance Workshop. Padua, Italy.
- Hirning M. 2014. The application of luminance mapping to discomfort glare: A modified glare index for green buildings [PhD dissertation]. Queensland University of Technology (QUT).
- Hubai Z. 2014. qdslrdashboard. [accessed 2019 Jan 17]. <https://dslrdashboard.info/>.
- Ihaka R, Gentleman R. 1996. R: A language for data analysis and graphics. *J Comput Graph Stat*. 5(3):299–314.
- Illuminating Engineering Society of North America (IESNA). 2000. The IESNA lighting handbook. New York (NY): Illuminating Engineering Society.
- Inanici M. 2006. Evaluation of high dynamic range photography as a luminance data acquisition system. *Light Res Technol*. 38(2):123–136.

- Inanici M. 2009. Introduction to high dynamic range photography. 8th International Radiance Workshop. Boston, MA.
- Inanici M. 2010. Evaluation of high dynamic range image-based sky models in lighting simulation. *LEUKOS*. 7(2):69–84.
- Inanici M, Galvin J. 2004. Evaluation of high dynamic range photography as a luminance mapping technique. Berkeley (CA): Lawrence Berkeley National Laboratory.
- International Electrotechnical Commission (IEC). 1999. Colour measurement and management in multimedia systems and equipment. Part 2-1: default RGB Colour Space - sRGB. Geneva (Switzerland): Swedish Institute for Standards (SIS).
- Jacobs A. 2007. High dynamic range imaging and its application in building research. *Adv Build Energy Res*. 1(1):177–202.
- Jacobs A. 2012. Glare measurement with hdr photography. 11th International Radiance Workshop. Copenhagen, Denmark. [10.1094/PDIS-11-11-0999-PDN](https://doi.org/10.1094/PDIS-11-11-0999-PDN)
- Jacobs A, Wilson M. 2007. Determining lens vignetting with hdr techniques. Paper presented at: Light'2007. Varna, Bulgaria. doi:[10.1094/PDIS-91-4-0467B](https://doi.org/10.1094/PDIS-91-4-0467B)
- Jakubiec JA, Inanici M, Van Den Wymelenberg K, Mahic A. 2016a. Improving the accuracy of measurements in daylight interior scenes using high dynamic range photography. Paper presented at: PLEA 2016. Los Angeles, CA.
- Jakubiec JA, Van Den Wymelenberg K, Inanici M, Mahic A. 2016b. Accurate measurement of daylight interior scenes using high dynamic range photography. Paper presented at: CIE 2016. Melbourne, Australia.
- Jones NL, Reinhard CF. 2017. Experimental validation of ray tracing as a means of image-based visual discomfort prediction. *Build Environ*. 113:131–150.
- Jung B, Inanici M. 2018. Measuring circadian lighting through high dynamic range photography. *Light Res Technol*. 51(5):742–763.
- Kent MG. 2016. Temporal effects in glare response [PhD dissertation]. University of Nottingham.
- Kerr D. 2008. The proper pivot point for panoramic photography. [accessed 2017 Jun 13]. http://dougkerr.net/Pumpkin/articles/Pivot_Point.pdf.
- Kim MH, Kautz J. 2008. Characterization for high dynamic range imaging. *Comp Graph Forum*. 27(2):691–697.
- Mantiuk R, Krawczyk G, Mantiuk R, Seidel H-P. 2007. High dynamic range imaging pipeline: perception-motivated representation of visual content. Paper presented at: SPIE: Human Vision and Electronic Imaging XII. San Jose, CA.
- Mitsunaga T, Nayar SK. 1999. Radiometric self calibration. Paper presented at: IEEE Conference on Computer Vision and Pattern Recognition. Fort Collins, CO.
- Pierson C, Piderit MB, Wienold J, Bodart M. 2017a. Discomfort glare from daylighting: influence of culture on discomfort glare perception. Paper presented at: CIE 2017 Midterm Meeting. Jeju, South Korea.
- Pierson C, Wienold J, Jacobs A, Bodart M. 2017b. Luminance maps from high dynamic range imaging: calibrations and adjustments for visual comfort assessment. Paper presented at: Lux Europa 2017. Ljubljana, Slovenia.
- Rea MS, Jeffrey I. 1990. A new luminance and image analysis system for lighting and vision I - Equipment and calibration. *LEUKOS*. 19(1):64–72.
- Reinhard E, Ward G, Pattanaik S, Debevec P. 2006. High dynamic range imaging - acquisition, display, and image-based lighting. San Francisco (CA): Elsevier Inc.
- Robertson MA, Borman S, Stevenson RL. 2003. Estimation-theoretic approach to dynamic range enhancement using multiple exposures. *J Electron Imaging*. 12(2):219–228.
- Stanley MJ. 2016. On the development and error analysis of a high dynamic range imaging system for luminance measurements [Master thesis]. Boulder (CO): University of Colorado Boulder.
- Stumpfel J, Tchou C, Jones A, Hawkins T, Wenger A, Debevec P. 2004. Direct hdr capture of the sun and sky. Paper presented at: AFRIGRAPH. Cape Town, South Africa.
- Van Den Wymelenberg K. 2012. Evaluating human visual preference and performance in an office environment using luminance-based metrics [PhD dissertation]. University of Washington (UW). [10.1094/PDIS-11-11-0999-PDN](https://doi.org/10.1094/PDIS-11-11-0999-PDN)
- Ward G. 1987. Radiance. [accessed 2018 Nov 28]. <https://radiance-online.org/>.
- Ward G. 1998. Findglare command (Radiance). [accessed 2017 Sept 27]. http://radsite.lbl.gov/radiance/man_html/findglare.1.html.
- Ward G. 1998. hdrngen. [accessed 2019 Nov 20]. <http://www.anywhere.com/>.
- Ward G. 2011. Raw2hdr perl script. [accessed 2019 Oct 25]. <http://www.anywhere.com/gward/pickup/raw2hdr.tgz>.
- Ward G. 2016. Fisheye_corr.Cal file. [accessed 2018 Nov 28]. https://www.radiance-online.org/cgi-bin/viewcvs.cgi/ray/src/cal/cal/fisheye_corr.cal.
- Wienold J. 2012. New features of evalglare. 11th International Radiance Workshop. Copenhagen, Denmark. [10.1094/PDIS-11-11-0999-PDN](https://doi.org/10.1094/PDIS-11-11-0999-PDN)
- Wienold J, Andersen M. 2016. Evalglare 2.0 – new features faster and more robust hdr-image evaluation. Paper presented at: 15th International Radiance Workshop. Padua, Italy.
- Wienold J, Christoffersen J. 2006. Evaluation methods and development of a new glare prediction model for daylight environments with the use of ccd cameras. *Energy Buildings*. 38(7):743–757.
- Wienold J, Reetz C, Kuhn TE, Christoffersen J. 2004. Evalglare - a new radiance-based tool to evaluate daylight glare in office spaces. Paper presented at: 3rd International Radiance Workshop. Fribourg, Switzerland.
- Wüller D, Gabele H. 2007. The usage of digital cameras as luminance meters. Paper presented at: Conference SPIE: Digital Photography III. San Jose, CA.