Research Article

Boris Lenseigne^{a,*}, Valéry Ann Jacobs^a, Martijn Withouck, Peter Hanselaer and Pieter P. Jonker

Color sensitivity of the multi-exposure HDR imaging process

Abstract: Multi-exposure high dynamic range(HDR) imaging builds HDR radiance maps by stitching together different views of a same scene with varying exposures. Practically, this process involves converting raw sensor data into low dynamic range (LDR) images, estimate the camera response curves, and use them in order to recover the irradiance for every pixel. During the export, applying white balance settings and image stitching, which both have an influence on the color balance in the final image. In this paper, we use a calibrated quasi-monochromatic light source, an integrating sphere, and a spectrograph in order to evaluate and compare the average spectral response of the image sensor. We finally draw some conclusion about the color consistency of HDR imaging and the additional steps necessary to use multi-exposure HDR imaging as a tool to measure the physical quantities such as radiance and luminance.

Keywords: color; HDR; spectral sensitivities.

^aBoth authors equally contributed to this work.

***Corresponding author: Boris Lenseigne**, Delft Biorobotics Laboratory, Delft, The Netherlands, e-mail: b.a.j.lenseigne@tudelft.nl
Valéry Ann Jacobs, Martijn Withouck and Peter Hanselaer: Light and Lighting Laboratory, KAHO Sint Leuven, Gent, Belgium
Pieter P. Jonker: Delft Biorobotics Laboratory, Delft, The Netherlands

1 Introduction

High dynamic range (HDR) imaging aims at producing images whose signal dynamics matches the variations of radiance in the scene. While a typical semiprofessional DSLR camera, like the Nikon D300 used in this experiment, has a typical dynamic range of 60 dB, the HDR images have a dynamic range far above 100 dB. The HDR images can be acquired using either specially designed sensors [1] or multiple exposures of the scene taken with an ordinary camera

with varying exposure times. Besides artistic photography, typical applications of the HDR images concern image-based rendering [2]. In most of these applications, the output of the imaging process is ultimately low dynamic range (LDR) data suitable for display or printing. Thus, most of the attention so far, concerning the colorimetric aspects of HDR images, has been focused on the color transcription onto a LDR display, a process named as tone mapping [3]. The HDR data, itself, is often described in a radiance map of the scene modulated by the spectral response of the imaging process. In their original work, Debevec and Malik [4] proposed a crude white balance step in their HDR stitching algorithm, while a solution has been proposed by [5] to produce HDR data with a predefined color profile, albeit not in the original RGB color space. Some other techniques have been proposed that operate directly on the HDR data (see for instance [6]), but the direct influence of the HDR image stitching itself is not well defined yet. This must be taken into account in order to fully understand how HDR imaging can be used to measure the physical quantities such as radiance or even the luminance of a scene, which is the radiance weighted by the visual response curve of a standard human observer.

This paper is a contribution to this study. We consider the case where the HDR images are built with multiple images of a static scene acquired from a fixed point of view with an ordinary DSLR camera. In this study, the scene consists of an integrating sphere, illuminated with a 50-W halogen lamp via a monochromator, which select a specific wavelength and, thus, produces a calibrated quasi monochromatic light (Cf. 2.1). Simultaneously, we measure the effective light intensity inside the sphere with a spectrometer, which enables us to recover the spectral response of each step imaging process (Cf. 2.4). We use 35 wavelengths from 380 to 720 nm, and for each of them nine successive exposures separated by one f-stop (Cf. 2.2). Next, we compare the spectral response of the raw sensor data over all the exposures with the one of the exported *tiff* images (Cf. 3.1), later used to build the HDR images (Cf.2.3). Finally, we draw some concluding remarks about the effects of each of these steps on the spectral response (Cf. 4).

www.degruyter.com/aot

2 Methods

2.1 Experimental setup

To calibrate the spectral response of the image sensor with respect to the HDR exposures, our setup creates a uniform emitting surface with a narrow spectral bandwidth by combining a light source, a monochromator, and an integrating sphere (Figure 1). The monochromator transmits a narrow bandwidth in the spectrum of a 50-W halogen. To reduce the harmonics that may arise at multiples of the selected wavelength, the light passes a filter wheel before entering the monochromator. Monochromatic light has only one color and is, thus, obtained by selecting only one wavelength. We selected light with a spectral bandwidth of 10 nm full width at half maximum (FWHM), which results in *quasi-monochromatic* light (Figure 2). The quasi-monochromatic light is coupled to the input port of a 6-inch integrating sphere of Edmund Optics, whose walls are quasi-Lambertian. In theory, a Lambertian surface reflects all light equally in every direction. A quasi-Lambertian surface approximates this reflectance behavior at best: our integrating sphere reflects and distributes the incoming light equally over the sphere's wall, which guarantees a uniformly emitting surface with a known spectral bandwidth.

A spectrograph is mounted on a telescopic lens to calibrate the camera's counts per second to the radiance in the red, green, and blue channel. The telescopic lens images the inside of the sphere – the uniform emitting surface – through an exit port. The focal length of the telescopic lens is chosen long enough so the field of view of the telescopic

lens is smaller than the exit port of the integrating sphere. The image of the telescopic lens is coupled to the spectrograph by an optical fiber so that for each selected bandwidth, a spectrum is gathered ranging from 229 to 903 nm. A calibration spectrum is taken of an irradiance standard (SRS8 – Bentham), which emits light homogeneously with a known spectral radiance. This known spectral radiance allows us to calibrate to the counts in the red, green, and blue channel. The absolute spectral radiance of the emitting disk can only be determined when calibrated to a known source. Therefore, we replaced the emitting disk by a calibration standard of spectral radiance (SRS8 – Bentham). The SRS8 is fitted with a 50-W halogen lamp that is connected to a constant current DC power supply at 4.000 A. Standard calibration range is from 380 to 800 nm, and it has a color temperature of 3260 K. To determine the spectral radiance for each wavelength-step, we calculate:

$$L_{e,\lambda,\text{test}} = \frac{\text{counts}_{\text{test}}}{t_{\text{test}}} \times \frac{W / m^2 \cdot sr \cdot nm}{\frac{\text{counts}_{\text{reference}}}{t_{\text{reference}}}} \quad (1)$$

where the second term in the product above is known for the calibration lamp. This procedure allows us to express the test source for each wavelength-step in units of spectral radiance ($W/(m^2 \cdot sr \cdot nm)$).

2.2 Multi-exposure images acquisition

We use a Nikon D300 DSLR camera (Nikon Corporation, Tokyo Japan) with an 85-mm AF Nikkor lens. The aperture is set to f/1.8 and ISO to 1000. The white balance is set

Figure 1 To determine the average spectral response of the image sensor, the setup we configured consists of a 50-W halogen, a filter wheel (FW) to reduce harmonics, a monochromator with a bandwidth of 10 nm, a 6-inch integrating sphere, a Nikon 300 camera, and an optical fiber that connects the telescopic lens to the spectrograph. To calibrate the setup by an irradiance standard, the field of view of the telescopic lens is chosen to be smaller than the exit port of the integrating sphere.

Figure 2 The calibrated spectrograph data showing 34 color bands as measured inside the integrated sphere. The 610-nm band has been removed because of measurement error.

using the same calibration light (SRS8 – Bentham) as the one used to calibrate the spectrometer. The exposure times are, respectively: 1/30, 1/15, 1/8, 1/4, 1/2, 1, 2, 4, and 8 s and are set so that the green channel (480 nm) would be under-exposed by one f-stop on the central exposure time (1/2s).

The images are saved at maximal resolution in the Nikon native *nef* format. The raw images are then converted to tagged image file format (*tiff*) with *ufraw* (<http://www.ufraw.sourceforge.net>), using camera white balance (color temperature equal to 3200 K) and adaptive homogeneity-directed (AHD) [7] interpolation for color recovery from the Bayer filter. The images used for single-exposure response calibration are exported with a 16-bit resolution, while the images used to build the HDR files are saved in an 8-bit format, with all other settings identical. The lower bit depth for the images devoted to HDR stitching is necessary because estimating the camera response requires solving the system for each possible pixel value, which is not possible directly for 16-bit values as the computing load becomes cumbersome.

2.3 Camera curve calibration and HDR stitching

The sensor response curves link every possible pixel value to certain irradiance. Most commonly, the curves are estimated separately for the red, green, and blue channels. In

this work, this calculation is performed using the implementation provided by Debevec and Malik [4].

Being given P successive exposures of the same static scene with known exposures times, $\Delta t_j (j=1..P)$, if we assume that all pixels at the same position i in the successive image have received the same irradiance, we can write the sensor response function as:

$$Z_{ij} = f(E_i \Delta t_j) \quad (2)$$

This is assumed to be a monotonic function; thus, by taking the logarithm of its inverse, we obtain:

$$g(Z_{ij}) = \ln(E_i) + \ln(\Delta t_j), g = \ln(f^{-1}) \quad (3)$$

The goal is to find, for every possible discrete pixel value, the appropriate value for $g(z)$ as well as every pixel radiance, E_i , from a set of N sampled locations in the image. As the solution can only be known up to a scale factor α , an additional constraint, enforcing the response function to be 0 for pixels with the value $Z_{mid} = (Z_{min} + Z_{max})/2$ is introduced as well as a weighting function, which emphasizes the pixels with a value closer to Z_{mid} to account for the lower smoothness of the response curve near its extremities:

$$w(z) = \begin{cases} z - Z_{min} & \text{if } z \leq Z_{mid} \\ Z_{max} - z & \text{if } z > Z_{mid} \end{cases} \quad (4)$$

In order to find a least square solution to this problem, we need then to minimize the following function, wrt, the

intensities E_i , and the camera curve values for each pixel value, $g(z)$:

$$\min_{E_i, g(z)} \left(\sum_{i=1}^N \sum_{j=1}^P [w(Z_{ij})(g(Z_{ij}) - \ln(E_i) - \ln(\Delta t_j))]^2 + \lambda \sum_{z=Z_{\min}+1}^{z=Z_{\max}-1} [w(z)g''(z)]^2 \right) \quad (5)$$

where the second term is a smoothness term that uses the discrete second derivative of the response function. This system can be solved directly using a singular value decomposition (SVD).

Once the response function, g , is known, knowing also the exposures, Δt_j , we can recover the HDR pixel level log radiance using equation (3) and the weighting function (4):

$$\ln(E_i) = \frac{\sum_{j=1}^P w(Z_{ij})(g(Z_{ij}) - \ln(\Delta t_j))}{\sum_{j=1}^P w(Z_{ij})} \quad (6)$$

In this experiment, the same curves are used to generate all the HDR images. Because the images made during our experiment have a limited spectral bandwidth and often exhibit saturation for the highest exposure, a separated sequence of nine exposures was made with the same camera settings to calculate the camera response curves. The calibration images consist of nine successive exposures of a Macbeth color calibration card onto a white wall, illuminated indirectly with a low-consumption light.

The light source is positioned so that the images would not saturate for the longest exposures. The resulting curves are shown in Figure 3. These curves are then used to generate 35 HDR images of the integrating sphere illuminated by 35 calibrated wavelengths and photographed under nine successive exposures.

2.4 Spectral response recovery

The procedure for the spectral response recovery is identical for all the raw, *tiff* and HDR images and will, thus, be described indifferently for either kind of image. It is similar to the one proposed by Sigernes et al. [8, 9].

For each image, a similar rectangular region of 200×200 pixels is cropped out. As, for very long exposures or extreme wavelength, the signal is corrupted by salt and pepper noise despite the dark frame removal, we first reduce that noise using a median filter of size 7×7 on each color band. For each image, the reference signal intensity is the average pixel value of the denoised patch and normalized w.r.t. the exposure time. The monochromator is tuned to the wavelength range 380–720 nm. By 35 steps of 10 nm, while the spectrograph provides us with the spectral radiant flux of the light, inside the integrating sphere, it provides 1024 steps in the 229–903 nm range. The calibrated spectral power distribution is in units of $W/(m^2 sr nm)$. The pixel values for each wavelength is then related to the spectral power distribution of the light by a sensitivity vector representing the exposure-compensated

Figure 3 The camera response curve showing the luminance captured by the pixels as a function of their 8-bit value. The curves are calculated by the Debevec and Malik algorithm.

response of the red, green, and blue channels, respectively, for each wavelength produced by the monochromator:

$$u_i = C \cdot \tilde{S}_i \cdot \Delta\lambda \text{ for } i=1..3 \quad (7)$$

with u_i as the vector of the normalized pixel values, C as the matrix of the spectrograph data subsampled to the wavelength produced by the monochromator, $\Delta\lambda$ as the band pass of the spectrograph, and $\tilde{S}_i$ as the vector of the unknown sensitivities for each color band. C is then a matrix containing 35 spectra in successive columns describing the spectral power distribution of light for each wavelength and 35 lines corresponding to the wavelength where the applied spectrum has been selected by the monochromator. Such a system can be solved directly for each color band ($i=1,2,3$) corresponding to the red, green, and blue channels in the image.

The unit in which the sensitivity is expressed depends on the type of data that is considered: for the raw image data, the normalized pixels values are in counts per second and the sensitivity in *CNT per W/(m².sr.nm.s)*. For the exported images, the sensitivity is expressed as *pixel value per W/(m².sr.nm.s)*. Finally, the HDR image is built up as to recover the radiance value for each pixel, and the sensitivity of the whole imaging process can now directly be expressed in *W/(m².sr.nm.s)*. However as, without prior calibration, the real radiance map can only be known up to a scale factor, the same applies to the spectral sensitivity.

3 Results

Several image conversion steps are necessary to build the HDR files: first, the images are acquired as raw sensor data, second, the raw data are converted into conventional images by color interpolation over the Bayer pattern and applying colorimetric settings (white balance). Finally, these images (their 8-bit counterparts in our case) are stitched together using the camera response curves calculated *a priori* (Cf. 2.3). We first describe how the spectral sensitivity varies over the different exposures for the raw and *tiff* image formats and, finally, the spectral sensitivity of the HDR images build from *tiff* image data.

3.1 Raw spectral response

The raw sensor data are extracted from the original *nef* files by first exporting them as uncompressed digital negative files (DNG) using the DNG Image Converter (<http://www.digikam.org>) (version 1.2.0). The images are then loaded

into MATLAB (Mathworks Inc., Natick, MA, USA), and the red, green, and blue color planes are separated. The spectral response is calculated as described in Section 2.4. The pixel counts are expressed in counts per second (*CNT per W/(m².sr.nm.s)*) and the wavelength in nanometers (nm). We first present the results for a single exposure: 1/8 s, which is the longest exposure that does not exhibit saturation over the whole range of wavelength. These values will be used as reference when studying the HDR spectral response (Cf. 3.2). The pixels counts (Figure 4A) in the blue channel reflect the fact that our light source's intensity is very low at wavelengths below 420 nm. This fact can also be observed on the spectrograph data (Figure 2). Similarly, one also observes the higher intensities around the 650-nm wavelength in the red channel counts. The corresponding spectrum (Figure 4B) exhibits a small predominance of the sensitivity of the green channel over the red and blue ones, which have almost similar peak values. The red peak is also slightly narrower than the blue one. The peak values for the red, green, and blue channels, respectively, are situated at 590, 540 and 470 nm. Surprisingly, the experiment conducted by Sigernes et al. [8] with a similar camera exhibited a different behavior where the blue channel would be predominant over the green and red ones, but these authors used different camera settings and especially a higher 1600 ISO sensitivity as well as 4 s of exposure time. The use of a different measurement unit also makes it difficult to compare the absolute sensitivity values. However, it is worth noticing that, in both cases, the curves have a similar shape, with especially a distinctive red peak at 530 nm and a smooth descent of the blue sensitivity up to 600 nm.

Next, we repeat the same experiment for the 16-bit exported *tiff* images. These images have been exported from the camera raw files using the camera white balance settings (Cf. 2.2). These are the images, in their 8-bit counterparts, that will be used to generate the HDR images. As, for now, the direct relationship between the radiance and the pixel value is lost (it will be recovered during the HDR generation process), we express the sensitivity as a *normalized pixel value per W/(m².sr.nm.s)* (normalization is w.r.t. to the exposure time). The normalized pixels values and pixel sensitivity, calculated as described in Section 2.4 for the exposure 1/8 s is shown in Figure 4C and D, respectively. It appears clearly, here, that the process of converting the images by applying specific white balance settings has enhanced the blue channel compared to the two others: in Figure 4C, the blue pixels appear to have almost a similar peak height as the green ones, whereas the raw pixels counts (Figure 4A) were much lower. Consequently, the corresponding sensitivity is boosted and appears now to be sensibly higher than the one of the other channels.

Figure 4 Exposure of 1/8 s, pixels counts, and sensitivity. Upper row: raw camera data. (A) Raw pixel counts, (B) raw sensitivity. Lower row: exported 16-bit *tiff* images. (C) Normalized pixels values. (D) Pixel sensitivity. Camera white balance.

The comparison with the raw spectrum above also shows that the conversion process has shifted the peak value of the green channel by 10 nm (from 540 to 530 nm). Together with the sharper peak of the red channel, this is the only significant change in the shape of the sensitivity curves that occurred during the conversion beside the adjustment of the respective gains of the color channels.

Over successive exposures, the raw sensitivity also appears to be very stable: Figure 5 shows the pixel counts (above) and the sensitivity (below) for the red, green, and blue channels and for the exposures 1/30 s (-4 eV), 1/8 s (-2 eV), 1/2 s (0 eV), 2 s (+2 eV), and 8 s (+4 eV). The main visible differences between the different exposures occur when the channel is saturated, which is also seen as a drop in the corresponding sensitivity as for the +2 and +4 eV curves. Besides that, the only other visible difference is a slight drop in the sensitivity at the peak value for the shortest exposures. This phenomenon appears in all the color

channels. For the shortest exposures, another phenomenon appears: an excessive sensitivity of all the channels below 400 nm. This appears to be images where the pixels counts are very low so that this phenomenon is very likely to be induced by the signal's background noise. The pixel sensitivity for the exported 16-bit images, however, exhibits less consistency (Figure 6). Besides the artifacts introduced by saturation for the highest exposures, which, once again, makes it difficult to draw conclusions for these exposures, it appears clearly that the pixel sensitivity for the exposure 0 eV (1/2 s) is almost half the one for the lowest exposures. For these low exposures (1/2, 1/8, and 1/30 s), the pixels' peak sensitivity drops as the exposure increases. The phenomenon is less visible in the green channel but appears clearly in the red one where the peak sensitivity at 1/2 s exposure is only one third of the one at 1/30 s. Things are even more dramatic in the blue channel where the sensitivity for the 1/8 s series also shows a more narrow peak than

Figure 5 (A) The camera raw data and sensitivity over successive exposures: raw pixel counts (above) and pixel sensitivity (below) for the red, green, and blue channel (from left to right) over different exposures. The graph shows the curves for the exposures 1/30 s (-4 eV), 1/8 s (-2 eV), 1/2 s (0 eV), 2 s (+2 eV), 8 s (+4 eV). (B) Close up view of the central part of the curves.

Figure 6 (A) Exported 16-bit Tiff images. Pixel values and sensitivity over successive exposures: raw pixel counts (above) and pixel sensitivity (below) for the red, green, and blue channels (from left to right) over different exposures. The graph shows the curves for the exposures 1/30 s (-4 eV), 1/8 s (-2 eV), 1/2 s (0 eV), 2 s (+2 eV), 8 s (+4 eV). Camera white balance. (B) Close up view on the central part of the curves.

the other exposures. In the blue channel also, the 1/30 s exposure curve shows the lowest sensitivity, while the peak sensitivity for the blue channel also shifts over the successive exposures: from 450 nm at 1/5 s to 480 nm at 1/8 s and 470 nm at 1/30 s. This shifting of the peak sensitivity is also present in the green channel (albeit less drastically) with peak values at 530 nm at 1/5 and 1/30 s exposures and at 540 nm at 1/5 s exposure time. This phenomenon, however, does not appear in the red channel.

These observations prove that, despite the fact that the camera sensor, itself, shows a well balanced spectral response and a good stability over the variable exposures, the process of exporting the raw camera data to another image format, which involves applying white balance settings, even if this format has the bit depth necessary to capture the whole dynamic range of the sensor, and the export was made with settings as neutral as could possibly be, has a huge influence on the spectral response of the imaging process. It was, of course, expected that the white balance settings would influence the respective balance of the red, green, and blue channels, but the variations that were observed in the pixel sensitivity for varying exposures also suggest that the export introduces additional nonlinearities in the dynamics of the imaging process.

3.2 Spectral response of the HDR imaging process

The last step concerns the influence of the HDR image stitching on the spectral response of the imaging process. The HDR images are built, for each wavelength, using the nine-exposure *tiff*-converted images downsampled to 8 bits, the algorithm described in Section 2.3, and a unique

camera curve (Figure 3) is used. As the HDR stitching intends to recover the radiance of the scene, we suppose that HDR pixel values are, up to an unknown scale factor, directly in units of radiance $W/(m^2.sr.nm)$. Keeping in mind the scale factor between HDR pixel value and actual radiance, the HDR sensitivity can be expressed directly in $W/(m^2.sr.nm.s)$, also up to an unknown scale factor.

The resulting pixel values, as shown in Figure 7, appear this time to give more emphasis to the red channel w.r.t. the green and blue ones when compared to the normalized pixel values for the exported *tiff* images (Figures 4C and 6): the red peak pixel value (at 600 nm) is now, indeed, almost twice as high as the green peak (at 540 nm), while this ratio was only 1.5 for the 1/8-s exposure image (Figure 4C). However, the peak values are located on the same frequencies (480, 540, and 600 nm for the blue, green, and red channels, respectively). As a result, the HDR spectral response shows its highest sensitivity in the red channel, while the green channel is now the one with the lowest maximal sensitivity. It is also interesting to observe that the HDR images exhibit the same increase in sensitivity, for all the bands, at a wavelength below 400 nm as we observed for the lower exposures in the raw images (Cf. 3.1) but not in the *tiff* images. This observation tends to prove that the camera response curves, estimated prior to the stitching, itself, indeed, somewhat mimics the behavior of the original sensor curves. However, even if the camera response is known, this result also shows that the radiance recovery is impaired by the spectral manipulations during the raw data conversion and also that the stitching, itself, influences the color balance of the images. Still, in this case at least, the spectral sensitivity of the HDR image is closer in shape to the raw spectrum (Figure 4B) than the *tiff* spectrum was.

Figure 7 HDR pixel values (A) and spectral sensitivity (B) for the nine exposures covering a range of eight f-stops.

These observations show that the process of forming an HDR image by stacking together multiple exposures of the same scene is not transparent toward the colorimetric settings of the source images.

4 Discussion

This study concerns the evolution of the spectral response of the different steps leading from raw pixel counts to the HDR radiance data. During this conversion, an intermediate step is necessary where the raw sensor data is converted into the conventional LDR image data prior to the multi-exposure fusion leading to the HDR radiance data. It appears that this step involves applying colorimetric transformations to the raw data (white balance) and also nonlinear dynamic manipulations even when the conversion settings were manually set to be as neutral as possible. As a result, the raw to LDR conversion modifies not only the balance between the color channels but also the shape of the response curves over varying exposures, so that the color consistency can actually not be guaranteed over the successive exposures, despite the fact that the original raw data exhibits a remarkably well-balanced spectral response, extremely stable over varying exposure times except signal saturation.

However, the HDR radiance recovery is based on the estimation of the camera response curves calculated over the LDR image data. These curves are, thus, expected to account for all the transformations that occurred previously, and we could, indeed, see that the estimated camera curves show a strong nonlinear behavior for extremely low pixel values. But it also appeared that the HDR data's spectral response differs from both the LDR and raw responses in terms of color balance, despite the fact that the color band shapes are similar to the raw ones up to a different scale factor for each color band. This result was expected, knowing that the HDR stitching estimates the radiances up to a scale factor and separately for each color band. On the other hand, this also shows that this effect cannot be neglected when one considers the colorimetric properties of the HDR data. This suggests that an additional calibration step is necessary in order

to produce data with a spectral response identical to the one of the sensor, which, at least for our camera, would be particularly interesting as the raw data showed a sensitivity close to the one of the standard human observer. This would make it possible to directly measure the luminance from the HDR radiance data, once the respective scaling factors for the red, green, and blue channel are known.

5 Concluding remarks

The main findings of this study can be summarized as follows:

1. We use an integrating sphere, illuminated by a quasi-monochromatic light, produced by a halogen lamp coupled to a monochromator to produce a diffuse area with a known spectral bandwidth.
2. We measure the light intensity for 35 different wavelengths with a spectrometer, attached to a telescopic lens and a Nikon D300 DSLR camera and match this data to evaluate the raw spectral response of the camera and the spectral response of the conventional image data over successive exposures.
3. Whereas the raw response is very stable over varying exposures, the exported data are not only modified by the white balance settings but also by the nonlinearities in the signal's dynamics introduced by the conversion.
4. While the camera response curves estimated prior to the HDR stitching account for these nonlinearities, unknown scale factors in the HDR radiance recovery still introduce unknown changes in the color balance of the final data w.r.t. the raw spectral one.
5. The sensor of our camera proved to be factory built to exhibit a sensitivity close to the one of the standard observer. Thus, building the HDR data with the original spectral sensitivity would make it possible to directly measure the luminance values from this data.
6. Estimating the radiance scale factors for each color band is, however, necessary so that either an improved stitching algorithm, or an additional calibration protocol, must be found.

Received January 6, 2013; accepted February 9, 2013

References

- [1] J. N. Burghartz, H.-G. Graf, C. Harendt, W. Klingler, H. Richter and M. Strobel, in 'Proceedings 8th International Conference on Solid-State and Integrated Circuit Technology' (2006).
- [2] E. Reinhard, G. Ward, S. Pattanaik and P. Debevec, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, ISBN: 0-12-585263-0(2005).

- [3] M. Čadík, M. Wimmer, L. Neumann and A. Artusi, J. Comput. Graphics 32-3 (2008).
- [4] P. E. Debevec and J. Malik, SIGGRAPH 97 (1997).
- [5] M. Goesele, W. Heidrich and H.-P. Seidel, in '9th Color Imaging Conference Color Science and Engineering: Systems, Technologies, Applications' (2001).
- [6] H. Yoo, K. Y. Kim and K. H. Lee, IEICE Trans. Inf. Syst. (2009).
- [7] K. Hirakawa and T. W. Parks, Trans. Image Process. 14-3 (2005).
- [8] F. Sigernes, M. Dyrland, N. Peters, D. A. Lorentzen, T. Svenøe, K. Heia, S. Chernouss, C. Sterling Deehr and M. Kosch, Opt. Express 17–22 (2009).
- [9] F. Sigernes, J. M. Holmes, M. Dyrland, D. A. Lorentzen, T. Svenoe, K. Heia, T. Aso, S. Chernouss and C. S. Deehr, Opt. Express 16–20 (2008).

Boris Lenseigne defended his PhD in Paul Sabatier University (Toulouse, France) in 2004 in sign language analysis and worked as a researcher for the Institut Pasteur Korea (Seoul, South Korea) where he was developing pattern recognition methods for confocal microscopy images analysis. He is, since 2008, a postdoctoral fellow in Vision-based Robotics at the Bio-Mechanical Engineering group of the Delft University of Technology (Netherlands). His current interests are about computer vision, image processing, cognitive models and perception, 3D reconstruction for humanoid robots, and alternative 3D reconstruction strategies including depth from motion and shape from shading.

Valery Ann Jacobs joined the faculty of engineering of the Vrije Universiteit Brussel (Belgium) as a PhD student in 2011, after obtaining a master in Astronomy and Astrophysics at K.U.Leuven. Her main research activity is related to non-imaging optics, but she has a keen interest in imaging optics as well.

Martijn Withoutack joined the engineering department of K.U.Leuven as a PhD student in 2011, after obtaining a master in Physics and Astrophysics at UGent. His main research activity is related to color perception.

Peter Hanselaer received his PhD in Physics at the University of Gent (Belgium) in 1986. Peter is professor at the Catholic University College Sint Lieven and associate professor at the K.U.Leuven, Dept. EAST/ELECTA (Elektrische Energie en Computer Architecturen). In September 2011, he became vice president of the department of industrial engineering, responsible for research and international cooperation. He is a board member of the Belgian Institute on Illumination and Belgian delegate in the CIE (Commission Internationale de l'Eclairage), division 1. In 1997, Peter founded the Light&Lighting Laboratory which was supported by IWT Flanders (agentschap voor Innovatie door Wetenschap en Technologie) and several industrial and scientific partners.

Pieter Jonker is a full professor in Vision-based Robotics at the Bio-Mechanical Engineering group of the Delft University of Technology (NL). He coordinated various multidisciplinary national and EU projects in the area of flexible manufacturing, robot assembly, real-time computer architectures, nano-computing, and autonomous cooperating robots. With Dr Martijn Wisse, he runs the Dutch Bio-robotics Laboratory at Delft University of Technology, in which bio-inspired research is performed on humanoid robots and robotics for care. He is a member of the IEEE and fellow of the IAPR. His current research interest is on bio-inspired real-time embedded vision systems for robotics, surveillance, and augmented reality and on hierarchical reinforcement learning for walking robots.