

Adaptive exposure estimation for high dynamic range imaging applied to natural scenes and daylight skies

Miguel A. Martínez,* Eva M. Valero, and Javier Hernández-Andrés

Color Imaging Laboratory, Department of Optics, Faculty of Sciences, University of Granada, Granada, Spain

*Corresponding author: m.martinez.domingo@gmx.com

Received 2 September 2014; revised 15 December 2014; accepted 17 December 2014;
posted 17 December 2014 (Doc. ID 222306); published 26 January 2015

Digital imaging of natural scenes and optical phenomena present on them (such as shadows, twilights, and crepuscular rays) can be a very challenging task because of the range spanned by the radiances impinging on the capture system. We propose a novel method for estimating the set of exposure times (bracketing set) needed to capture the full dynamic range of a scene with high dynamic range (HDR) content. The proposed method is adaptive to scene content and to any camera response and configuration, and it works on-line since the exposure times are estimated as the capturing process is ongoing. Besides, it requires no *a priori* information about scene content or radiance values. The resulting bracketing sets are minimal in the default method settings, but the user can set a tolerance for the maximum percentage of pixel population that is underexposed or saturated, which allows for a higher number of shots if a better signal-to-noise ratio (SNR) in the HDR scene is desired. This method is based on the use of the camera response function that is needed for building the HDR radiance map by stitching together several differently exposed low dynamic range images of the scene. The use of HDR imaging techniques converts our digital camera into a tool for measuring the relative radiance outgoing from each point of the scene, and for each color channel. This is important for accurate characterization of optical phenomena present in the atmosphere while not suffering any loss of information due to its HDR. We have compared our method with the most similar one developed so far [IEEE Trans. Image Process. **17**, 1864 (2008)]. Results of the experiments carried out for 30 natural scenes show that our proposed method equals or outperforms the previously developed best approach, with less shots and shorter exposure times, thereby asserting the advantage of being adaptive to scene content for exposure time estimation. As we can also tune the balance between capturing time and the SNR in our method, we have compared its SNR performance against that of Barakat's method as well as against a ground-truth HDR image of maximum SNR. Results confirm the success of the proposed method in exploiting its tunability to achieve the desired balance of total Δt and SNR. © 2015 Optical Society of America

OCIS codes: (110.1758) Computational imaging; (100.2000) Digital image processing; (040.1490) Cameras; (010.7295) Visibility and imaging; (110.1085) Adaptive imaging; (110.0110) Imaging systems.
<http://dx.doi.org/10.1364/AO.54.00B241>

1. Introduction

Natural scenes are usually composed by a wide variety of radiance signals outgoing from the objects in the scene, which are very different in magnitude. This fact makes their correct capture with a normal digital camera a nontrivial problem. Capturing a scene with

high dynamic range (HDR) content with a single low dynamic range (LDR) image would cause loss of information in those regions of the scene where the light level reaching the sensor is too low or too high to be correctly registered with a single exposure time. The HDR imaging techniques solve this problem.

Common imaging sensors suffer from limitations in the process of capturing the light. Usually, the dynamic range of the sensor (i.e., the ratio between the maximum and minimum irradiance impinging

on the sensor that produces an effective response) is much lower than the dynamic range found in natural open air scenes. The dynamic range of these scenes (ratio between the maximum and minimum radiances emitted by the objects in the scene) can vary from 2 to 8 orders of magnitude depending on the season and scene content [1]. The human visual system can simultaneously adjust to a difference of up to 3.73 orders of magnitude (or log units) [2] when adaptation is accomplished. However, most imaging and display devices can only account for barely 2 orders of magnitude in a single image (either for capturing or for displaying) [3,4].

There have been many techniques proposed [5–7] as well as sensor architectures [8,9] to achieve this goal. The most common techniques are based on building a HDR image from the information of a number of different LDR images. The difference between these LDR images is the exposure (i.e., the product of irradiance impinging in the sensor times the exposure time used to acquire the image). It is changed by varying either exposure time or the aperture. Usually, the former is adjusted since it does not affect the depth of field between different captures. These differently exposed versions of the same scene, when combined in the correct way, can be used to build an image that contains extended dynamic range information compared to just a single exposure [10].

When we use our camera to capture a digital image of a scene, we cannot know in advance which exposure times would be useful for composing the HDR image afterwards. We could just take a large number of images with different exposure times (or even all available ones in the camera), and then use all of these LDR images to compose the HDR image. But this option is often time consuming and very computationally demanding, so it is not always feasible. Commercial cameras usually have an auto-exposure mode that estimates an exposure time value based on some cues like the reading from a built-in photometer in high-end cameras that measures the average brightness value in certain regions of the image.

All of these cases aim just for finding one exposure time that works well for correctly imaging most parts of the scene. However, whether we would be able to find it or not, there is not a single value of exposure time that could make all pixels in one shot be correctly exposed for most common natural scenes.

The aim of this paper is not to explain how to merge LDR images into a HDR radiance map. The process we used for this is very well explained in the literature [11]. It is rather to present a method for the selection of a set of exposure times (bracketing set) to use in order to retrieve useful information from all pixels (or at least from most of them). This is very important for the study of optical phenomena present in the atmosphere and open air natural scenes with shadows, twilights, clouds, crepuscular rays, and so forth [12], all of which have HDR content that cannot be captured with a single shot of a commercial digital camera. Thus, a digital camera can be

a useful tool for composing a HDR radiance map of these phenomena in order to study them. Of course, if the scene captured has very dark regions that need long exposure times to be correctly exposed, then it is important that there is no relative movement between the camera and scene content during the capturing time. If small movements happen, there are ghosting-compensation techniques to correct for artifacts [13].

We aimed for a method that is blind (no information from the scene content is known *a priori*), adaptive (adapts to scene content dynamically by adjusting required exposure times), universal (works for any camera that we have tested so far), and on-line (the exposure times are calculated as the capturing process is ongoing and every single shot acquired is used in the HDR radiance map generation). It will also give as default output the minimal bracketing set (the bracketing set that has the minimum number of shots, yet recovers the full dynamic range of the scene), but it can be tuned to yield longer exposure times with a higher signal-to-noise ratio (SNR). This tunability is introduced as a method that controls the amount of overlapping between consecutive exposures to increase the SNR in the resulting HDR image at the cost of increasing the number of shots taken, and hence the capturing time. We have also introduced a method to control the percentage of pixel population that we can accept to be useless.

The remainder of this paper is organized as follows. Section 2 summarizes the state-of-the-art technology for estimating exposure times for HDR imaging. Section 3 explains the details of the method we propose. Section 4 explains the experiments made to compare our method with the most advanced method of those described in Section 2, as well as the results obtained. Finally, Section 5 draws the main conclusions of this work.

2. State of the Art

Several approaches have been proposed in the literature for solving the problem of finding the exposure time values for HDR image capture via multiple exposures. Chen and Gamal [14] proposed scheduling for capture times. They were assuming a known illuminant in the scene, which in practice is a rather non-realistic assumption, especially for optical open air phenomena. Grossberg and Nayar [15] proposed a method to simulate the response of any camera (linear, logarithmic, gamma function, etc.) using a single camera with a known camera response function (CRF) by just selecting a set of exposure times. So their aim was not to find a minimum bracketing set for radiance map generation. Stumpf *et al.* [16] proposed a method for capturing HDR images of the sun and sky. They threshold the images and check if there are saturated and/or underexposed pixels. If any, they add new shots by increasing or decreasing the exposure by a fixed amount of three-stops. This approach is not adaptive to the scene content and could lead to situations where the exposure times

are not well fitted to cover the full dynamic range of the scene. Bilcu *et al.* [17] proposed a method for overcoming the limitations of mobile devices for HDR imaging. Their proposal was also done by iteratively trying every available value of exposure time and afterwards deciding which is the correct one. Therefore, many images need to be taken before a bracketing set is selected.

Barakat *et al.* [18] proposed a method for finding minimal bracketing sets for HDR capture. Firstly, they studied how the camera responds to radiance using every available exposure time. Then, they select only those exposure times that completely cover the full dynamic range of the camera with certain overlap. This is the so-called minimal system bracketing set (MSBS), and whatever the content of the scene being imaged is, using all these exposures will always cover the full radiance range that the camera can effectively acquire. To adapt this MSBS to scene content, they proposed to select a subset of it called the minimum image bracketing set (MIBS) by capturing a first shot with an intermediate exposure time (that belongs to the MSBS) and checking if there are saturated or underexposed pixels. If so, they add the next exposure time included in the MSBS until the full dynamic range is covered. Though the underlying idea in this method is similar to the one proposed in this work, it is still not totally adaptive to scene content since they limit the exposure times selected to those belonging to the MSBS. We believe that the same scenes could be captured with less shots and shorter exposure times, yet covering their full dynamic range, and we have demonstrated this by comparing the results of our method with Barakat *et al.*'s insightful algorithm using their MIBS approach.

Granados *et al.* [19] proposed a method assuming the known mean HDR irradiance histogram of the scene being captured. Besides, their method only works for linear cameras, and they used a greedy algorithm, iteratively capturing the same scene many times until they obtained the optimal SNR solution. Thus, their method is not on-line. Hirakawa and Wolfe [20] used a mathematical method based on training for HDR exposure time selection. They assumed linear sensor response and known noise sources in the capture, which is not always a realistic scenario. They did not really aim for defining minimum bracketing sets, but for optimal SNRs instead. Gelfand *et al.* [21] adapted HDR imaging to mobile devices as well. They merge LDR images iteratively two by two. If there are still saturated or underexposed pixels, they keep adding the next available exposure time that the camera offers. Hasinoff *et al.* [22] proposed a method to calculate a bracketing set that is optimal in terms of the SNR but not minimal (they just try to fit it within a given time budget) by varying both the exposure time and the ISO (i.e., sensitivity) settings of the camera. Besides, their method assumes linear raw sensor responses only and known information about scene radiance content.

Gallo *et al.* [23] proposed a method for taking advantage of mobile phone camera APIs (application programming interfaces). They programmed the automatic histogram calculation in mobile phones to construct a reduced HDR histogram of the scene, which will be the target to be captured. This is however not possible if the camera used for the captures does not feature this automatic process. Besides, if the scene imaged has very dark regions, the long exposures needed to create this histogram make the process slow. Moreover, the method is not on-line. After this histogram is calculated, they capture several exposures of the scene and then study many possible combinations of them until the optimal one is selected.

Guthier *et al.* [24] followed the lines previously proposed by Kang *et al.* [25], who implemented a method for sequentially adjusting the exposure for real-time HDR video. Both are iterative and limited to only two shots for building the HDR image. Finally, Gupta *et al.* [26] proposed a Fibonacci-series-based bracketing set determination algorithm in which each exposure time is the sum of the previous two. This technique does not aim for full dynamic range recovery though, but image registration for HDR video.

3. Proposed Method

We drove our Canon EOS 7D camera from our laptop via the USB (universal serial bus) port using the open-source libraries called *GPhoto2* from our algorithm implemented in Matlab R2014a and working on-line.

The method proposed in this paper is full range because it finds a bracketing set that covers the full dynamic range radiance map of the scene. This HDR radiance map would be potentially useful for studying the behavior of light in HDR open air phenomena.

Our method uses the CRF to compute the relative irradiance impinging on the sensor, which corresponds to a certain population of pixels in the image (using its cumulative histogram). Then, a new exposure time is calculated in order to shift the camera responses to this irradiance to a different value. This way, the cumulative histogram is shifted and the new shot would capture a different range of irradiances, which are contiguous to the range captured in the previous shot.

The CRF is a function that relates the response of the camera, in digital counts (DC), with the exposure that the sensor receives. This function depends on each camera, and it even can be different for different settings of the same camera (e.g., a camera working in raw mode or in jpeg mode). Knowing the CRF of our camera is a key factor to build the radiance map. A detailed explanation of how to calculate the CRF is given in [11]. The calculated CRF of the camera used for the experimental part in this work is shown in Fig. 1. It is clearly not linear.

The exposure axis is in relative units, and it is normalized so that the center of the DC values (128 DC for the 8-bits case) corresponds to a relative exposure

value of 1. The function is the same for the three color channels R, G, and B of the camera, since it is a property of the sensor. Therefore, we process the three color channels together like the technique in [23]. For each LDR image we capture, we know the Δt used as well as the DC values for each pixel and channel. Therefore, by using the CRF we can easily work out the relative irradiance (E) by computing the simple ratio shown in Eq. (1):

$$E_{i,k,e} = \frac{H(i,k,e)}{\Delta t(e)}. \quad (1)$$

The subindex i accounts for the pixel index, k accounts for the color channel, and e accounts for the exposure index (or number of shots). Thus, once we have captured an initial image with a known exposure time Δt_0 , the CRF relates H_0 with DC_0 as shown in Eq. (2):

$$DC_0 = \text{CRF}(H_0) = \text{CRF}(E_0 \cdot \Delta t_0). \quad (2)$$

Therefore, we can work out the relative irradiance value of a point of the image [E_0 , calculated as shown in Eq. (3)] by knowing the CRF, the Δt , and its DC value in the first shot (Δt_0 and DC_0 , specifically):

$$E_0 = \frac{\text{CRF}^{-1}(DC_0)}{\Delta t_0}, \quad (3)$$

where CRF^{-1} refers to the inverse CRF function that always exists since CRF is a monotonically increasing function. Then, to shift the sensor responses DC_0 to this same irradiance value E_0 into a new value DC_1 , we just have to work out which new exposure time Δt_1 is needed for a new shot, like shown in Eq. (4):

$$\Delta t_1 = \frac{\text{CRF}^{-1}(DC_1)}{E_0}. \quad (4)$$

If our camera has only a limited set of values from which to choose the exposure time, we can select the available value that is closest to the calculated one.

We already have a tool to control the values of sensor responses, which is done by tuning the exposure time used to acquire the images. Now, we explain how to use it for our purpose of optimizing HDR capture. For this aim, we propose a method based on cumulative histograms of the scene inspired by

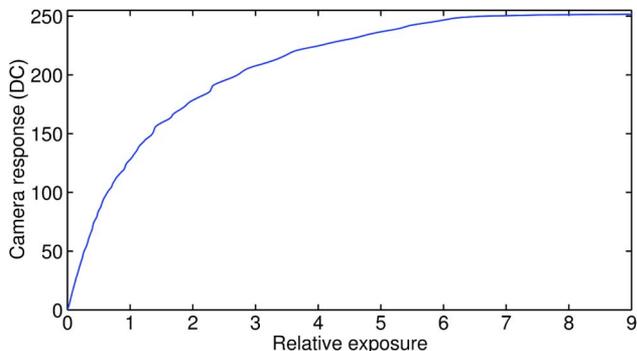


Fig. 1. CRF of the Canon EOS 7D camera in jpeg mode.

Grossberg and Nayar [27], who originally applied it to pixel selection for CRF computations. If the scene content does not change, then the same value for the percentile of population in the cumulative histogram will correspond to the same areas in the image. In Fig. 2, we can see a plot where several cumulative histograms of the same scene differently exposed are drawn together.

If a given percentile is below some exposure value for a given exposure time, then, for a different exposure time, the same percentile of population will correspond to a different exposure value but they will still keep its location within the scene. Therefore, the points where the horizontal lines in Fig. 2 intersect the histograms report information corresponding to the same areas of the scene. Our idea is to shift sensor responses by calculating exposure times to control the sensor responses to pixel populations of key percentile values. As a starting point, we calculate the cumulative histogram of the image captured with the automated exposure of the camera. But in principle, any image can be used as starting point as long as it has some pixels that are neither saturated nor underexposed.

We are going to sample the scene's radiance using the CRF of the camera between two DC levels. Unless the scene has a very reduced dynamic range, there will be pixel values below and above these DC values. Since in the default version of the algorithm we aim for minimum bracketing sets, we have set the low level (LO) to 3 DC and the high level (HI) to 252 DC for considering a pixel to be underexposed or saturated, respectively, when it is out of these bounds. Thus, whatever pixel population is above the HI level or below the LO level, we will sample it using a different exposure time. Here, we introduce two novel features of our method. One is the possibility of setting a tolerance for the percentage of useless pixels. If we choose 0% tolerance, the algorithm will look for longer or shorter exposure times if at least one pixel is underexposed or saturated, like Barakat *et al.* [18] proposed. However, for some scenes we can set a different tolerance threshold to renounce to a certain percentage of the population to be properly exposed [e.g., when we directly image the sun and our region of interest (ROI) is in a different area].

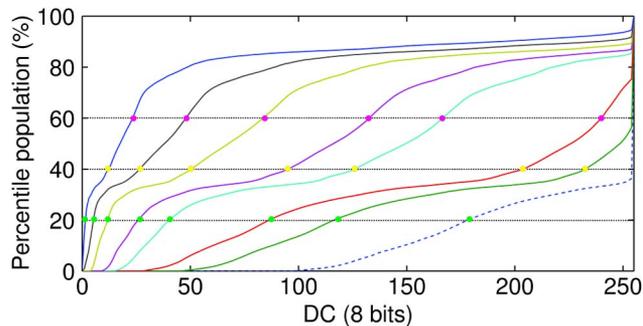


Fig. 2. Cumulative histograms of the same scene captured using different exposure times.

The other novelty involves controlling the LO and HI values of the CRF. Setting values very close to the extremes (0 and 255 DC for 8 bits) will result in a lower number of shots at the cost of a lower SNR. In contrast, if we set values further from these extremes, we will sample the scene's radiance with more overlap between contiguous shots and therefore the SNR will increase, at the cost of a higher number of shots. This shows how our algorithm can be tuned to adapt to different requirements regarding the SNR of the captured HDR.

After commenting on these functionalities, we describe now in detail how the exposure time search is done. With the information present in the cumulative histogram of the first shot captured, we check the percentile of the pixel population that is below the LO level. If it is higher than the maximum value set, then a longer exposure time is calculated. The same is done if the difference between 100 and the percentile of the pixel population above the HI level is higher than the tolerance threshold. In this case, a shorter exposure time will be calculated.

To find a longer exposure time, we will use Eq. (4) to shift the camera response value from the LO level to HI level. Therefore, we use the HI level as DC_1 and E_0 is substituted by Eq. (3), where we use the LO level as DC_0 . Δt_0 is the exposure time used to acquire the current image [see Eq. (5)]:

$$\Delta t_{\text{longer}} = \frac{\text{CRF}^{-1}(\text{HI})}{\text{CRF}^{-1}(\text{LO})} \cdot \Delta t_0. \quad (5)$$

In contrast, to find a shorter exposure time, we will use Eq. (4) to shift the camera response value from

the HI level to LO level. Therefore, we use the LO level as DC_1 and E_0 is substituted by Eq. (3), where we use the HI level as DC_0 [see Eq. (6)]:

$$\Delta t_{\text{shorter}} = \frac{\text{CRF}^{-1}(\text{LO})}{\text{CRF}^{-1}(\text{HI})} \cdot \Delta t_0. \quad (6)$$

In this way, if the population that has a sensor response at the HI level in one shot shifts to the LO level in the next shot, we can cover the full dynamic range of the scene with certain overlap between contiguous shots.

The process described here goes on checking the cumulative histograms of the longer and shorter exposure times until the tolerance requested is met or the system reaches its maximum or minimum available exposure times.

4. Experiments and Results

Our camera (Canon EOS 7D) allows the choice of only a discrete set of exposure times. We tuned the HI and LO levels (explained in Section 3) to get the minimum bracketing sets (lowest SNR). In the first experiment, we tested the default version of the adaptive exposure estimation (AEE) method (see Section 4.B.1); in the second experiment, we explored the tunability and evaluated the SNR performance (see Section 4.B.2). In all scenes tested, the method built a full dynamic range radiance map of the scenes. We implemented our proposed method as well as the only method found following the same philosophy, which is the MIBS method proposed by Barakat *et al.* [18] (hereafter termed BAR), to compare their performances.

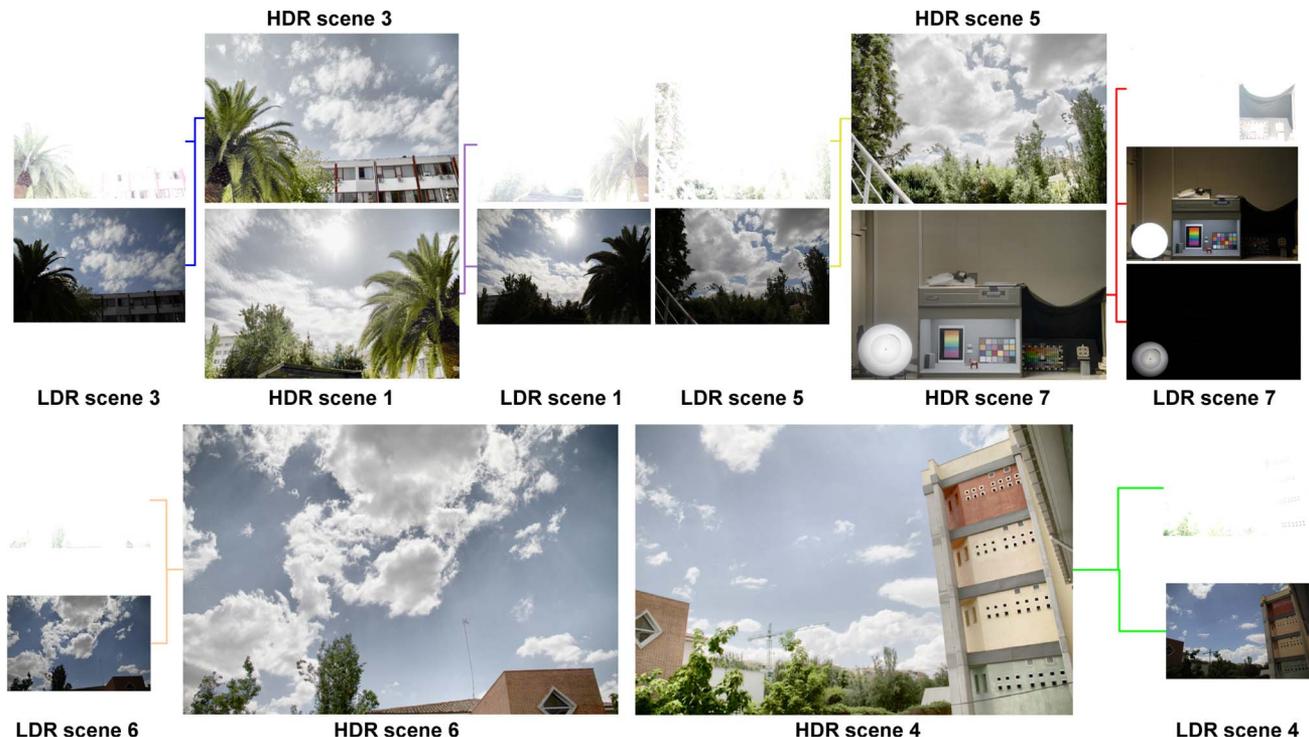


Fig. 3. LDR images and HDR radiance maps for some captured scenes.

A. Bar Method and MSBS

Regarding the BAR method, the MSBS found for our Canon camera, as explained in [18], using a 5.6 aperture setting, was composed by four exposure times for which the values were: 30 s, 300 ms, 1 ms, and 0.0125 ms. Sometimes, not all of these four shots were needed to record the full dynamic range of the scene. In these cases, a sub-set of the MSBS is used omitting some of its shots. This represents the MIBS. We could use as well as a first shot, the one chosen by the auto-exposure mode of the camera, as was done for the AEE method. However, we found that when doing so, we only got the same number of shots or even one more. So we did not use it. We see an example of this in Fig. 4. The exposure time chosen by the auto-exposure mode of the camera was 66.7 ms (top-center). It was used as first shot for the AEE method, since this method adapts to any exposure value chosen as first. However, this value was in between 300 and 1 ms (both belong to MSBS). Therefore, if we used it also as a first shot for the BAR method, it would mean that the capture of this scene would end up with five shots instead of four.

Also, if we change the aperture setting of the camera, as the CRF is not changing, the AEE method would work the same by adapting to the new exposure levels impinging in the sensor. However, for the BAR method we would need to calculate a new MSBS, since the same exposure times for a different aperture would not be valid any more. Thus, we fixed our aperture setting to 5.6 for both methods.

B. Comparison between the AEE and BAR Methods

For the first experiment (Section 4.B.1), we captured 30 scenes using both methods and we studied the number of shots taken, the total exposure time used, and the percentage of the pixel population that was not properly exposed. In this way, we assessed how efficiently did both methods recover the full dynamic range of the scene by comparing their resulting bracketing sets.

In the second experiment (Section 4.B.2), we built an indoors HDR scene with controlled illumination conditions. We captured 10 HDR images of it using the AEE and BAR methods. Besides, for the AEE method we repeated the capture four times using different values for the LO and HI levels (see Section 3). Finally, we captured 10 ground-truth (GT) HDR images using all available exposure times in the camera. These GT images represent the highest SNR that our camera can achieve to record a scene without repeating shots with the same exposure.

1. HDR Capturing Efficiency

As mentioned before, we acquired 30 HDR scenes; 23 scenes were captured outdoors with natural illumination and 7 were captured indoors with artificial illumination. Outdoors, daylight cast HDR illumination over objects including clouds. Indoors we used a light booth and a fluorescent lamp oriented directly to the camera in a dark room to generate HDR content. To check the performance of both methods in terms of full range recovery, we plotted the

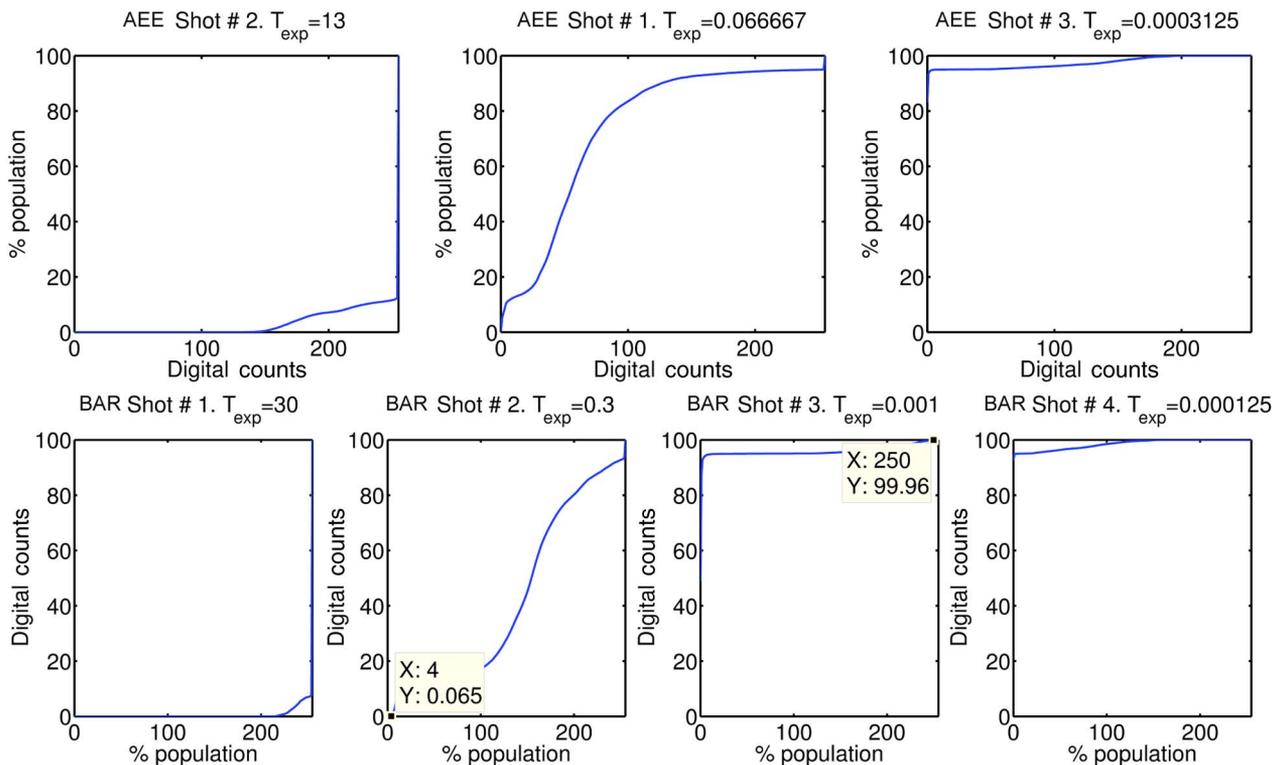


Fig. 4. Cumulative histograms of the same scene using the AEE method (top row) and BAR method (bottom row). The histograms are ordered by decreasing exposure time to observe their continuity.

cumulative histograms of all shots taken for each scene and checked that no irradiance gaps were left uncovered between consecutive shots. We set the maximum percentile allowed to be lost to 0%. In Table 1, we can see the results for 7 out of the 30 scenes captured.

We observed how the number of shots is always equal or lower for the AEE method. The percentage of useless pixels is always the same for both methods. Many scenes had a percentage of lost pixels equal to 0, since both methods managed to retrieve the full dynamic range of the scene. For the rest of the scenes, the useless pixels were due to direct sunlight (like the case of scene 1). This made some pixels impossible to recover even for the shortest exposure time available in the camera. The total exposure time is always lower for the AEE method. The same trends commented on were found for the remaining 23 scenes captured. In total, for the 30 scenes captured, the BAR method took a total of 96 shots using 218.734 seconds and the AEE method took a total of 81 shots using 139.869 seconds. This means that the number of shots was 15.63% less, and the exposure time was 36.06% less, for the AEE method.

In Fig. 3, you can see the LDR pictures and the tone-mapped HDR radiance maps generated for some of the scenes. The tone-mapping algorithm used was a contrast-limited adaptive histogram equalization, which was introduced by Ward [28] and implemented in Matlab R2014a.

In Fig. 4, we plot an example of the cumulative histograms corresponding to both methods for scene number 7. The data-tips in histograms for shots 2 and 3 for the BAR method highlight that there are still underexposed and saturated pixels in those exposures, although at first glance the histograms may seem to reach percentiles 0 and 100, respectively.

We observed how both methods succeeded in recovering the full dynamic range of the scene. However, thanks to the adaptation of the AEE method, only three shots were needed instead of the four shots the BAR method used. Therefore, we can conclude with this experiment that the AEE method recovers the

dynamic range of the scene as well as the BAR method does, but by using a more reduced bracketing set.

2. Signal-to-Noise Ratio

We did a second experiment to study the SNR behavior of our method. For the AEE method, we tested four different conditions named A (LO = 3 and HI = 253), B (LO = 16 and HI = 240), C (LO = 56 and HI = 200), and D (LO = 106 and HI = 150). We compared all these AEE conditions with the BAR method and the ground-truth (GT) images.

The dynamic range of the scene was measured using a spectroradiometer (Photo Research, PR-745) to measure the integrated radiances of both the brightest and the darkest points of the scene. The resulting dynamic range measured was 4.1 log units.

For each pixel and each color channel of these HDR radiance maps, we calculated its average HDR value and its standard deviation across the ten images corresponding to each method. The average HDR value provides information about the signal level in the pixel, and the standard deviation provides information about the level of noise generated by all noise processes present in the HDR capture process. Thus, by computing the ratio of the average HDR value (\bar{E}_{xy}) over the standard deviation (σ_{xy}) as Eq. (7) shows, we obtain a SNR estimate [29]:

$$\text{SNR}_{xy} = 20 \times \log_{10} \left(\frac{\bar{E}_{xy}}{\sigma_{xy}} \right). \quad (7)$$

The subindex xy stands for pixel position within the HDR radiance map.

We can see the number of shots, the total exposure times, and the average SNR for each method in Table 2.

As expected, setting the LO and HI values further from the extremes of the range in the AEE method yields a higher number of shots and also higher total exposure time, but the SNR increases as well. For condition D, we reached an average SNR only less than 2 dB below the ideal case (GT), yet using only about 40% of the total exposure time. The minimum bracketing set found was AEE A, with only three shots and 14.79 s of total exposure time, but this also had the lowest SNR. The BAR method needed the full MSBS to recover this scene using four shots. It had a better SNR than our minimum bracketing

Table 1. Minimum Bracketing Sets Found for BAR and AEE Methods in 7 of the 30 Scenes Captured

Scene #	Method	# Shots	% Lost	$\sum_{i=1}^n T_i (s)$
1	BAR	3	1.61	0.301
	AEE	2	1.61	0.025
2	BAR	3	0	0.301
	AEE	2	0	0.101
3	BAR	3	0	0.301
	AEE	2	0	0.040
4	BAR	3	0	0.301
	AEE	2	0	0.050
5	BAR	3	0	0.301
	AEE	2	0	0.034
6	BAR	3	0	0.301
	AEE	2	0	0.025
7	BAR	4	0	30.301
	AEE	3	0	13.067

Table 2. SNR Performance for Four AEE Conditions and the GT and BAR Methods

Method	# Shots	Δt (s)	$\overline{\text{SNR}}$ (dB)
AEE A	3	14.79	27.19
AEE B	4	15.02	30.47
AEE C	5	31.32	32.58
AEE D	16	61.22	33.57
GT	55	151.43	35.32
BAR	4	30.30	29.97

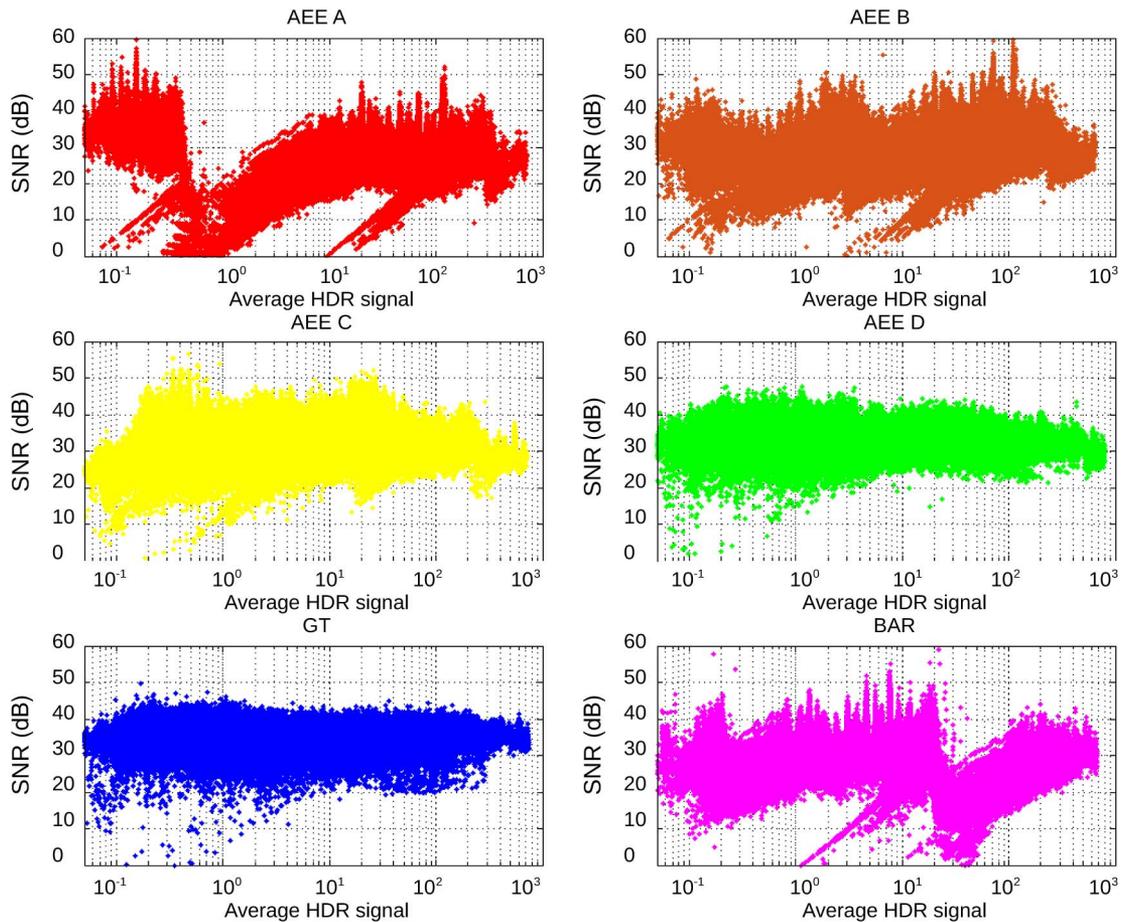


Fig. 5. SNR versus average HDR signals present in the radiance maps.

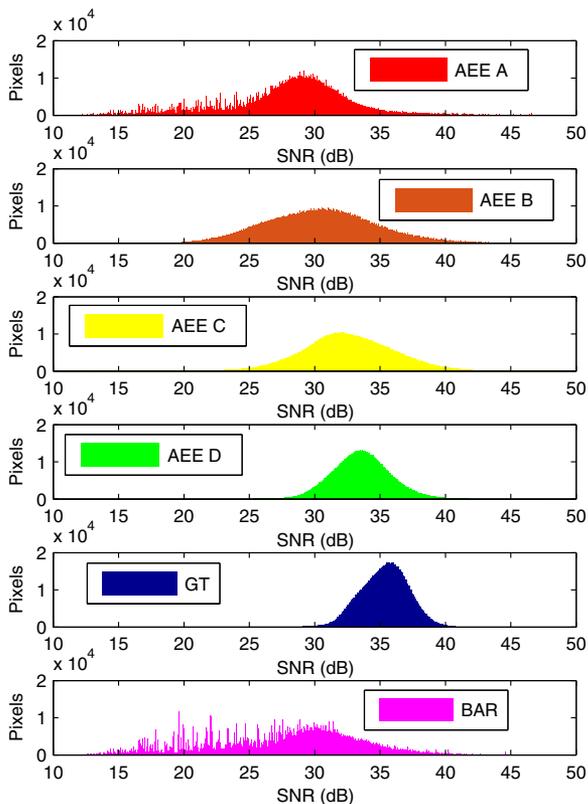


Fig. 6. Histograms of SNR values.

set, but we also had a second option (AEE B) using four shots with a shorter total exposure time (less than half) and a higher SNR than the BAR method.

We can observe in Fig. 5 the SNR for each pixel of the radiance maps generated versus the signal level. We can also see how the AEE D has the most similar distribution compared with GT. Also, the BAR method has a very similar distribution compared with cases AEE A and AEE B, as expected.

Figure 6 plots the SNR histograms for all methods. We can observe how for the AEE method, the main lobe gets narrower and shifts towards a higher mean SNR as we tune the LO and HI levels further from the extremes of the range. The AEE D is quite close to GT in position and shape. In contrast, the BAR method yields a histogram that is the mostly spread over a wide range of SNR values.

5. Conclusions

We present a new method for estimating the exposure times needed to recover the HDR radiance map from a scene via multiple exposures. We compared the performance of our method with that of the only method found in literature that aims for the same purpose (i.e., finding minimum bracketing sets) and performs under the same conditions (adapted to scene content with no *a priori* information about it and valid for any camera whether it is linear or not).

Our proposed method is adaptive because it finds a bracketing set adapted to any HDR scene content, and it is universal because it works for any camera. We only need to calculate its CRF (which is needed anyway to build the radiance map).

Moreover, the method is tunable, since we can decide if we prefer to find a minimum bracketing set at the cost of higher SNR or increase the SNR sampling the radiance of the scene with more overlapping between consecutive shots (increasing the number of shots and capturing time as well). For the minimum bracketing set case, the bracketing sets found were minimal in the 30 scenes tested.

Futhermore, our method is blind, which means that no information about the content of the scene needs to be known *a priori*. The multiple LDR images are captured on-line as the process is ongoing, and every single shot taken is used to compose the HDR radiance map.

We can also control the percentage of the total pixel population that we can assume is useless (underexposed or saturated). This way, we can find the minimum bracketing set only for our region of interest.

We have applied the method for HDR imaging of natural scenes where partially cloudy skies were present in order to increase the dynamic range of the capture. We have successfully covered the full dynamic range of the 30 scenes imaged. We have shown how our method can find bracketing sets that are shorter than those found by the BAR method, yet keeping higher SNR levels in the HDR radiance map reconstructed from the multiple exposures.

We studied the SNR performance of our method comparing it not only with the BAR method but also with an ideal-case ground-truth HDR image built using all available exposure times in the camera. We have demonstrated how we can tune our method to suit different requirements for the SNR at the cost of increasing the number of shots.

The proposed method brings a solution for the blind acquisition of HDR images using multiple exposures, which can be used in any HDR imaging context: machine vision, sky imaging, daylight illuminated scenes, HDR photography, etc. And in particular, the proposed method may be useful for studying optical phenomena present in open air scenes where the illumination conditions are extreme (i.e., direct sunlight that might be surrounded by regions of interest like halos, clouds casting shadows, and rainbows, just to give a few examples).

This work was funded by the Spanish Ministry of Economy and Competitiveness through the research project DPI2011-23202.

References

1. J. J. McCann and A. Rizzi, *The Art and Science of HDR Imaging* (Wiley, 2011), Vol. 26.
2. T. Kunkel and E. Reinhard, "A reassessment of the simultaneous dynamic range of the human visual system," in *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization* (ACM, 2010), pp. 17–24.
3. K. Yamada, T. Nakano, S. Yamamoto, E. Akutsu, and I. Aoki, "Wide dynamic range vision sensor for vehicles," in *Proceedings of Vehicle Navigation and Information Systems Conference* (IEEE, 1994), pp. 405–408.
4. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," in *ACM Transactions on Graphics* (ACM, 2004), Vol. 23, pp. 760–768.
5. S. Mann and R. Picard, *Being Undigital with Digital Cameras* (MIT, 1994).
6. S. Nayar and T. Mitsunaga, "High dynamic range imaging: Spatially varying pixel exposures," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2000), pp. 472–479.
7. S. Nayar and V. Branzoi, "Adaptive dynamic range imaging: Optical control of pixel exposures over space and time," in *Proceedings of the Ninth IEEE International Conference on Computer Vision* (IEEE, 2003), pp. 1168–1175.
8. S. Kavusi, K. Ghosh, and A. El Gamal, *Architectures for High Dynamic Range, High Speed Image Sensor Readout Circuits* (Springer, 2008).
9. S. Kavusi and A. El Gamal, "Quantitative study of high-dynamic-range image sensor architectures," in *Electronic Imaging* (International Society for Optics and Photonics, 2004), pp. 264–275.
10. M. A. Robertson, S. Borman, and R. L. Stevenson, "Dynamic range improvement through multiple exposures," in *Proceedings of the International Conference on Image Processing* (IEEE, 1999), pp. 159–163.
11. P. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *ACM SIGGRAPH 2008 Classes* (ACM, 2008), p. 31.
12. D. K. Lynch and W. C. Livingston, *Color and Light in Nature* (Cambridge University, 2001).
13. O. Gallo, N. Gelfand, W. Chen, M. Tico, and K. Pulli, "Artifact-free high dynamic range imaging," in *IEEE International Conference on Computational Photography (ICCP)* (IEEE, 2009), pp. 1–7.
14. T. Chen and A. El Gamal, "Optimal scheduling of capture times in a multiple-capture imaging system," in *Electronic Imaging* (SPIE, 2002), pp. 288–296.
15. M. D. Grossberg and S. K. Nayar, "High dynamic range from multiple images: which exposures to combine," in *Proceedings of the ICCV Workshop on Color and Photometric Methods in Computer Vision (CPMCV)* (IEEE, 2003).
16. J. Stumpfel, C. Tchou, A. Jones, T. Hawkins, A. Wenger, and P. Debevec, "Direct HDR capture of the sun and sky," in *Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa* (ACM, 2004), pp. 145–149.
17. R. C. Bilcu, A. Burian, A. Knuutila, and M. Vehvilainen, "High dynamic range imaging on mobile devices," in *the 15th IEEE International Conference on Electronics, Circuits and Systems, ICECS* (IEEE, 2008), pp. 1312–1315.
18. N. Barakat, A. N. Hone, and T. E. Darcie, "Minimal-bracketing sets for high-dynamic-range image capture," *IEEE Trans. Image Process.* 17, 1864–1875 (2008).
19. M. Granados, B. Ajdin, M. Wand, and C. Theobalt, "Optimal HDR reconstruction with linear digital cameras," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR* (IEEE, 2010), pp. 215–222.
20. K. Hirakawa and P. J. Wolfe, "Optimal exposure control for high dynamic range imaging," in *17th IEEE International Conference on Image Processing (ICIP)* (IEEE, 2010), pp. 3137–3140.
21. N. Gelfand, A. Adams, S. H. Park, and K. Pulli, "Multi-exposure imaging on mobile devices," in *Proceedings of the International Conference on Multimedia* (ACM, 2010), pp. 823–826.
22. S. W. Hasinoff, F. Durand, and W. T. Freeman, "Noise-optimal capture for high dynamic range photography," in *IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2010), pp. 553–560.
23. O. Gallo, M. Tico, R. Manduchi, N. Gelfand, and K. Pulli, "Metering for exposure stacks," *Computer Graphics Forum* 31, 479–488 (2012).

24. B. Guthier, S. Kopf, and W. Effelsberg, "Optimal shutter speed sequences for real-time HDR video," in *IEEE International Conference on Imaging Systems and Techniques* (IEEE, 2012), pp. 303–308.
25. S. Kang, M. Uyttendaele, S. Winder, and R. Szeliski, "High dynamic range video," in *ACM Transactions on Graphics (TOG)* (ACM, 2003), Vol. **22**, pp. 319–325.
26. M. Gupta, D. Iso, and S. K. Nayar, "Fibonacci exposure bracketing for high dynamic range imaging," in *IEEE International Conference on Computer Vision* (IEEE, 2013), pp. 1473–1480.
27. M. D. Grossberg and S. K. Nayar, "What can be known about the radiometric response from images?" in *Computer Vision ECCV* (Springer, 2002), pp. 189–205.
28. E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting* (Morgan Kaufmann, 2010).
29. J. T. Bushberg and J. M. Boone, *The Essential Physics of Medical Imaging* (Lippincott Williams and Wilkins, 2011).