HAL
open science

Best algorithms for HDR image generation. A study of performance bounds

Cecilia Aguerrebere, Julie Delon, Yann Gousseau, Pablo Musé

To cite this version:

Cecilia Aguerrebere, Julie Delon, Yann Gousseau, Pablo Musé. Best algorithms for HDR image generation. A study of performance bounds. 2012. hal-00733853v1

HAL Id: hal-00733853

https://hal.science/hal-00733853v1

Preprint submitted on 19 Sep 2012 (v1), last revised 6 Jun 2013 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

x
Vol. xx, pp. x
© xxxx
x–x

Best algorithms for HDR image generation
A study of performance bounds

Cecilia Aguerrebere[†‡], Julie Delon [†], Yann Gousseau [†], and Pablo Musé [‡]

Abstract. Since the seminal work of Mann and Picard in 1994, the standard way to build high dynamic range (HDR) images from regular cameras has been to combine a reduced number of photographs captured with different exposure times. The algorithms proposed in the literature differ in the strategy used to combine these frames. Many experimental studies comparing their performances have been reported, showing in particular that a maximum likelihood estimation yields the best results in terms of mean squared error. However, no theoretical study aiming at establishing the performance limits of the HDR estimation problem has been conducted. Another common aspect of all HDR estimation approaches is that they discard saturated values. In this paper, we address these two issues. More precisely, we derive theoretical bounds for the HDR estimation problem, and we show that, even with a small number of photographs, the maximum likelihood estimator (MLE) performs extremely close to these bounds. As a second contribution, we propose a general strategy to integrate the information provided by saturated pixels in the estimation process, hence improving the estimation results. Finally, we analyze the sensitivity of the HDR estimation process to camera parameters, and we show that small errors in the camera calibration process may severely degrade the estimation results.

Key words. high dynamic range imaging, irradiance estimation, exposure bracketing, multi-exposure fusion, camera acquisition model, censored data, exposure saturation, Cramér-Rao lower bound.

1. Introduction. The human eye has the ability to capture scenes of very high dynamic range, retaining details in both dark and bright regions. This is not the case for current standard digital cameras. Indeed, the limited capacity of the sensor cells makes it impossible to record the irradiance reaching very bright regions for long exposures. Pixels saturate incurring in information loss under the form of censored data. On the other hand, if the exposure time is reduced in order to avoid saturation, very few photons will be captured in the dark regions and the result will be masked by the acquisition noise. Therefore the result of a single shot picture of a high dynamic range scene, taken with a regular digital camera, contains pixels which are either overexposed or too noisy.

High dynamic range imaging (HDR for short) is the field of imaging that seeks to accurately capture and represent scenes with as large as possible irradiance range. The representation problem of how to display a HDR image or irradiance map in a lower range image (for computer monitors or photographic prints) while retaining localized contrast, known as tone mapping, will not be addressed here. Due to technological and physical limitations of current optical sensors, nowadays the most common way to reach high irradiance dynamic ranges is by combining multiple low dynamic range photographs, acquired with different exposure times $\tau_1, \tau_2, \ldots, \tau_T$. Indeed, for a given irradiance C and exposure time τ_i, the corresponding pixel value is a function of the received luminous energy $\tau_i C$. Hence, using different exposure times

[‡]Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay. pmuse@fing.edu.uy

[†]LTCI CNRS & Telecom ParisTech, 75634 PARIS Cedex 13, France. aguerreb@telecom-paristech.fr julie.delon@telecom-paristech.fr yann.gousseau@telecom-paristech.fr

allows to sample the camera response function at different operating points, thus avoiding saturation for at least some of the exposures and keeping details in both dark and bright regions. This acquisition process is called exposure bracketing, and is an automatic function in most modern digital SLR cameras.

For a given number of photographs, the choice of exposure times and the combination strategy are the "degrees of freedom" of HDR methods. Of course, the more photographs spanning the whole range of exposure times, the better. Actually ground-truth HDR images are built in such conditions, in a very controlled environment. However, for practical limitations, exposure bracketing in real situations can rarely exceed a few snapshots.

In this work we will concentrate on the problem of estimating the irradiance map, that is the irradiance reaching each pixel, from a reduced number of photographs captured with a given set of exposure times. More precisely, the general aim of this work is to establish the performance bounds for this estimation problem, and to find out how far from these bounds current state-of-the-art HDR estimation methods are. Thus, we will work under the hypothesis that the photographs are perfectly co-registered and possible radiometric changes have been compensated.

Previous work. To our knowledge, the first HDR imaging technique in the framework of digital photography, based on exposure bracketing, was proposed in 1994 by Mann and Picard [13]. The method assumes that the camera radiometric response function f mapping luminous energy to pixel values has been previously calibrated. Then, if z_i^p denotes the image value at pixel p for exposure time τ_i, the irradiance estimate at pixel p is computed as the weighted average of $f^{-1}(z_i^p)/\tau_i$, $i = 1, \ldots, T$. Small weights are assigned to extreme pixel values – very low or close to saturation – based on the claim that the camera response function calibration is less accurate at these values. This approach suffers from mainly two problems. The first one is that the weights are somehow arbitrary and not derived from a noise model of the pixel values. The second one is that the calibration of the camera response function is a problem on its own, and is prone to errors that are directly transferred to the irradiance estimator. The work of Mann and Picard represents an important contribution, since it inspired several approaches based on exposure bracketing, whose main difference relies in the way photographs are combined. See [8, 11] for an interesting review and comparison of these methods.

Luminance estimation methods can be classified according to different criteria. One of them is whether the method assumes a linear [9, 8, 11] or a non-linear camera response function [13, 2, 16, 21, 4].The former are meant to be used with the camera raw data, i.e. the pixel values before any camera post-processing (demosaicking, white balance, etc). For raw data, the camera response function is linear since, ignoring noise sources, each pixel value is proportional to the number of photons reaching the corresponding sensor cell. The latter needs to define a method to estimate the camera response function and its inverse. In this work we will assume a linear camera response. With the currently available technology and storage capacity of cameras, it is entirely reasonable to assume that we have access to the raw data or that the processing can be done directly on the raw data inside the camera. For those methods like Mann and Picard's, which consider non-linear response functions, we will therefore consider the linear response counterpart based on raw data.

Another meaningful classification of irradiance estimation methods is based on the ob-

jective function that is optimized. Three main groups can be distinguished, whether the objective function is based on the camera response function [13, 2], the SNR [16, 4] or the variance [21, 9, 8, 11, 19]. The methods in the first group [13, 2] propose to compute the weights based on the uncertainty of the sample values given by the camera response function. They claim that extreme pixel values are less accurate. Therefore they assign lower weights to those values and prioritize mid-range values. The SNR based methods [16, 4] weigh samples according to their SNR, which is computed from the input samples.

Variance based approaches are among the most recent ones. An idea common to all of them is the use of a statistical model of the camera acquisition process. This model takes into account several noise sources allowing to improve the irradiance estimation. Tsin et al. [21] were the first ones to propose this kind of approach. They characterize each pixel as a random variable whose distribution parameters depend on the unknown irradiance. Then they estimate the irradiance as an average weighted by the standard deviation of the samples. Robertson et al. [9] propose a statistical model where the mean pixel value depends on the unknown irradiance. Then they compute the irradiance with a maximum likelihood estimator (MLE). The MLE is the average of the input samples weighted by the inverse of their variances. However, since they do not manage to estimate these variances they fix their values by an ad-hoc procedure.

More recently, a new variance-based approach was introduced by Kirk and Andersen [8]. The main difference with the previous ones relies in the camera acquisition model. In this more realistic model both the mean and the variance depend on the unknown irradiance (because of the Poisson nature of the photon shot noise). They propose to estimate the irradiance also using a MLE. Now under this model a closed-form does not exist for the MLE. To overcome this limitation they assume a known variance at each pixel and they compute it using the pixel values. The same line of work is followed by Granados et al. [11] which propose a still more accurate model and solve the variance irradiance-dependence problem with an iterative method. Basically, at each iteration the weights are computed from the current irradiance value and the irradiance is updated according to the MLE (the one obtained ignoring the variance dependence on the irradiance).

Contributions. In this work we conduct a thorough analysis of the HDR image estimation problem, in order to establish its performance limits and to determine if current estimation techniques are close to these limits. Part of this study also consists in quantifying how these limits are affected when samples saturate, and when uncertainties are introduced in the camera parameters calibration process. We thus propose in this paper three interrelated contributions:

1. We present a study of the theoretical performance bounds of the HDR estimation problem. More precisely, given a small number of samples per pixel (say 6), we determine the lowest mean squared error (MSE) that can be attained by combining these samples. One question we address is why the MLE outperforms the other estimators proposed in the literature, and how far it is from the optimal MSE. The optimality of the MLE is far from obvious in such a non asymptotic case. We show however that there is not much room for improvement.
2. Surprisingly, all the methods proposed in the literature discard saturated samples. Nevertheless, saturated samples certainly carry some useful information. For instance the exposure time at which they have saturated. As a second contribution, we study

the usefulness of this information, and we present a way to incorporate it in the whole estimation process. Not surprisingly, the information contained in saturated samples turns out to improve the irradiance estimation.

3. The third contribution of this paper deals with a very significant problem that does not seem to have received much attention in the literature: the sensitivity of irradiance estimation to uncertainties in the camera parameters. This is a crucial question, since these parameters are obtained through a calibration process, which is of course not error-free. We show that irradiance can only be accurately estimated if the camera parameters are very carefully calibrated. Moreover, a very careful calibration is also critical in order to obtain a reliable ground-truth for evaluations with real data.

The article is organized as follows. In §2 a statistical model for the camera acquisition process is introduced. Then, in §3, based on the previous model, we derive performance bounds for the HDR estimation problem, and we compare some irradiance estimators, in particular the MLE, against the bound. The use of saturated samples in the irradiance estimation problem is analyzed in §4. In §5 we present a sensitivity analysis of the estimation problem with respect to uncertainties in camera parameters. The main conclusions of this work are summarized in §6.

2. Camera Model.

In this section, we briefly introduce the camera model that will be used throughout the paper. A more complete description of this model can be found in [1]. Two technologies are used for camera sensors: charge-coupled devices (CCD) and complementary metal-oxide-semiconductors (CMOS). A very similar acquisition model can be proposed for both sensors, which transform incoming light photons into voltage output values. We divide the main uncertainty sources of the acquisition process in two categories: random noise sources, and spatial non-uniformity sources.

2.1. Random noise sources.

Two physical phenomena are responsible for the random noise generation during the camera acquisition process: the discrete nature of light, which is responsible for photon shot noise, and the thermal generation of electrons.

Photon shot noise. The number of photons $\mathbf{C}_i^p$ impinging the photo-diode p during a given exposure time τ_i follows a Poisson distribution, with expected value $C_p\tau_i$, where C_p is the number of photons reaching p per unit of time. The number of electrons photoelectrically generated is also Poisson distributed and the voltage measured at the sensor output should be $V = g_{cv}C_p\tau_i$, where g_{cv} is the equivalent capacitance of the photo-diode.

Dark current. Some of the electrons accumulated on the potential well of the photo-diode result from thermal generation. The number of these electrons is well modeled by a Poisson distribution with expected value D^p [20], depending on the temperature and exposure time. In this paper, we call $\mathbf{D}_i^p$ this *dark shot noise* for the exposure time τ_i.

Readout noise. Some thermal noise affects the output values of the readout circuitry. This *reset noise* $\mathbf{N}_{reset}$ is accurately modeled as Gaussian distributed [12]. Another thermal noise $\mathbf{N}_{out}$, also modeled as Gaussian, appears during the amplification of the voltage values.

2.2. Spatial non-uniformity sources.

Besides random noise sources, several uncertainty factors, all related to the spatial non-uniformity of the sensor, should be taken into account in the acquisition model.

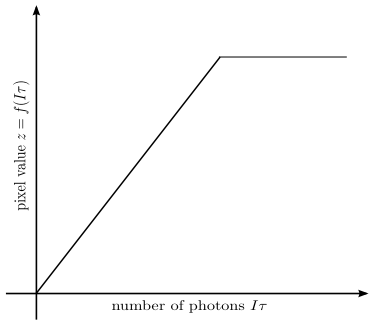

Figure 2.1: Camera response function for the linear + saturation model (raw data).

Photo-response non-uniformity (PRNU). Different pixels do not produce the same number of electrons from the same number of impacting photons. The fact that a photon can be absorbed or not in the photo-diode is a binomial selection of the Poisson process of impinging photons. Hence the PRNU can be modeled as a spatially variable multiplicative factor a_p applied to the parameter of the Poisson variable $\mathbf{C}_i^p$.

Dark-current non-uniformity (DCNU). The DCNU represents the variations in dark current generation rates from pixel to pixel. The DCNU can also be modeled as a spatially variable multiplicative factor d_p applied to the parameter of the Poisson variable $\mathbf{D}_i^p$.

2.3. Acquisition model. Equation (2.1) proposes a simplified model including the previous noise sources. Pixels are modeled as independent and following the same model (the dependence on position p is avoided to simplify the notation):

$$\mathbf{Z}_i = f\left([g_{cv}(\mathbf{C}_i + \mathbf{D}_i) + \mathbf{N}_{reset}]g_{out} + \mathbf{N}_{out}\right) + \mathbf{Q}, \tag{2.1}$$

where f is the camera response function, g_{out} is the gain of the camera, $\mathbf{Z}_i$ is the raw pixel value and $\mathbf{Q}$ is the uniformly distributed quantization error, which takes place during the conversion of the analog voltage measures into digital quantized values. The term between the round brackets in (2.1) can be rewritten as the addition of a Poisson distributed random variable with expected value $\lambda_i = a\tau_i C + dD_i$, multiplied by the gain factor $g = g_{cv}g_{out}$, and a Gaussian distributed noise component $\mathbf{N}_R = g_{out}\mathbf{N}_{reset} + \mathbf{N}_{out}$ with mean μ_R and variance σ_R^2.

Regarding the relative importance of each noise source, various articles agree in concluding that under low illumination conditions, the primary noise source is the reset noise, while for high illumination the major noise source is the photon shot noise [6, 5]. For the values normally taken by $\lambda_i = a\tau_i C + dD_i$, the Poisson distribution can be correctly approximated by a Gaussian distribution with mean and variance equal to λ_i. Moreover, the dark current can be neglected for exposure times below 1s [14] and it is commonly admitted that the quantization noise is negligible compared to the readout noise [6, 5]. As a consequence, we can assume that the variable $f^{-1}(\mathbf{Z}_i)$ follows a Gaussian distribution $N(ga\tau_i C + \mu_R, g^2 a\tau_i C + \sigma_R^2)$. In the case of raw data, f is a linear function before attaining its saturation threshold (see Figure 2.1).

Thus for non saturated samples the model becomes

$$\mathbf{Z}_i \sim N(ga\tau_i C + \mu_R, g^2 a\tau_i C + \sigma_R^2). \tag{2.2}$$

Since the camera parameters are assumed to be known, we will actually consider the $-\mu_R$ translated versions of these variables, and we will consider that their mean are linear in C.

3. Performance Bounds. This section is devoted to the irradiance estimation problem and to the computation of the corresponding performance bounds under the previous simplified acquisition model (2.2). Let us recall the basics of the problem. For each pixel of the sensor, we observe $\mathbf{z} = \mathrm{z}_1, \ldots, \mathrm{z}_T$ values for different exposure times $\tau_1, \ldots, \tau_T$ (each z_i is a realization of a random variable $\mathbf{Z}_i$ following the Model (2.2)) and we wish to estimate the irradiance C at this pixel, *i.e* the average number of photons reaching the pixel per unit time. This estimation will be done on a per-pixel basis, assuming that images are perfectly registered. Throughout this section, we assume that saturated samples have been discarded, as is done in all existent irradiance estimation methods. The inclusion of saturated samples and its impact on the estimators performances will be studied in Section 4.

After an introduction on the irradiance estimation as a statistical problem, the section addresses the computation of the Cramér-Rao lower bound for this model and the question of its attainability. We then proceed with the study of the comparative performance of existent estimators against the CRLB. We focus in particular on the MLE since it was found to outperform experimentally other methods in [8, 11].

3.1. Irradiance estimation: a statistical problem. Finding the optimal estimator $\hat{C}$ of C from the observations z_i is not obvious in practice. By "optimal", we mean that the estimator $\hat{C}$ should minimize the quadratic risk (or mean square error) $\mathrm{E}[|\hat{C} - C|^2]$. The main difficulties in this estimation come from the small sample size (between 2 and 6 samples in practice), and from the low SNR of the samples when C is low.

Biased or not biased?. The mean squared error (MSE) of $\hat{C}$ can be classically decomposed as the sum of its squared bias and its variance. One of the first questions that should be raised is thus the one of the right balance between these two terms. As underlined in the introduction, most of the irradiance estimators proposed in the literature consist of linear combinations of the values $\frac{\mathrm{z}_i}{ga\tau_i}$, with weights w_i summing to 1. In most cases, the weights depend on the input samples, thus the estimators are *a priori* biased. However, since $\mathrm{E}\left[\frac{Z_i}{ga\tau_i}\right] = C$, they are unbiased to a first order approximation [1]. Only two methods differ from the others in this regard: the approach of Debevec *et al.* [2], which is also unbiased to a first order approximation, but only as an estimator of $\log C$, and the MLE-based approach of Granados *et al.* [11], which is unbiased only asymptotically. In practice, we checked that for all the tested estimators, the ratio $\frac{\mathrm{E}[|\hat{C}-C|]}{C}$ remains very close to 0 whatever the value of C (see Section 3.4). This suggests that for all these estimators, the dominant term in the MSE is the variance. As a consequence, in the following we mainly focus on unbiased (or with negligible biases) estimators and on variance minimization.

[1] The delta method can be used to compute a first order approximation of the expected value of the estimators.

Is MLE *a good estimator of C?.* The second question coming naturally to mind concerning $\hat{C}$ is the one of the optimality of the MLE. Indeed, the MLE is known to be asymptotically efficient. However, in the HDR image generation problem the sample size is too small (normally in the order of 2 to 6) to consider the asymptotic approximation. Yet, even with a finite number of samples (less than 6 in our case), it can be shown that the MLE would be an optimal solution in terms of MSE if the variance $g^2aC\tau_i + \sigma_R^2$ in Model (2.2) were a constant or were proportional to C. The constant variance approximation is valid in very low irradiance conditions, resulting in a linear MLE, more precisely, it is the weighted average of the irradiance estimations for each exposure. On the other hand, the variance proportional to C approximation is valid in high irradiance conditions. The MLE for that case is not linear but has a closed-form. Now, under Model (2.2), the MLE does not have a closed-form. The solutions proposed in the literature to compute it numerically consist in simplifying the model [8] and making use of an iterative approach [11]. Nevertheless, the quality of the estimation in this general case cannot *a priori* be easily stated. One of the goals of this section is to answer this question.

3.2. Cramér-Rao lower bound for irradiance estimation. The previous paragraphs motivate the study of the performance bounds of the estimation problem. Once the expected performance limits are known, we may determine whether they can be reached, and compare the results obtained by existing estimation methods against them. This comparison states how close existing methods are to the limits, and allows to quantify how much room is left for improvement. The performance bound of the problem is given by the Cramér-Rao lower bound theorem. The Cramér-Rao lower bound states the minimum variance that we can expect to achieve (for a given bias). Knowing that the samples distribution is given by (2.2), the Cramér-Rao lower bound for unbiased estimators can be computed as (c.f. appendix A)

$$\text{CRLB} = \left[\sum_{i=1}^{T} \frac{(ga\tau_i)^2}{g^2a\tau_iC + \sigma_R^2} + \frac{(g^2a\tau_i)^2}{2(g^2a\tau_iC + \sigma_R^2)^2}\right]^{-1}, \tag{3.1}$$

where T is the number of considered exposure times. The MSE of any unbiased estimator is bounded by this expression. The immediate question to be raised is whether an efficient estimator exists (*i.e.* an unbiased estimator that attains the CRLB). The Cramér-Rao lower bound theorem states that an unbiased estimator may be found that attains the bound for all C if and only if the first derivative of the log-likelihood function can be factorized as

$$\frac{\partial \ln p(\mathbf{z}, c)}{\partial c} = I(c)(h(\mathbf{z}) - c), \tag{3.2}$$

for some functions h and I. In that case, the efficient estimator is $\hat{C} = h(\mathrm{z})$, and its variance is $1/I(C)$ [7]. Equation (3.2) must be valid for any c and any $\mathbf{z}$. Now, let us consider $\mathbf{z}$ fixed and take the limit of the two terms of (3.2) when c tends to infinity

$$\lim_{c\to\infty} \frac{\partial \ln p(\mathbf{z}, c)}{\partial c} = -\frac{1}{2}\sum_{i=1}^{T} \tau_i, \tag{3.3}$$

$$\lim_{c\to\infty} I(c)(g(\mathbf{z}) - c) = -\sum_{i=1}^{T} \tau_i. \tag{3.4}$$

The results in (3.3) and (3.4) are different, which proves that an efficient estimator does not exist for the problem.

Once the performance bounds of the problem have been determined, we are interested in the comparative performance of existing irradiance estimation methods against the CRLB. As previously stated, we focus on the MLE since it was found to outperform experimentally other methods in [8, 11]. The MLE for the irradiance under Model (2.2) does not have a closed-form. As a consequence, it is not obvious to directly compute its variance to evaluate its performance against the CRLB. Instead, we propose in §3.4 a detailed experimental study of its performance, together with that of other widely known irradiance estimation methods, relatively to this bound. Before the experimental study we present in §3.3 some theoretical hints of the optimality of the MLE. This result is verified experimentally afterwards in §3.4.

3.3. How close is the MLE from the CRLB: Some theoretical hints. The lack of closed-form for the MLE under Model (2.2) motivates Granados et al. [11] to propose an iterative algorithm that converges to the MLE solution. At iteration $(j+1)$ the irradiance $\hat{C}^{(j+1)}$ is estimated according to[2]

$$\hat{C}^{(j+1)} = \frac{\sum_{i=1}^{T} w_i(\hat{C}^{(j)})x_i}{\sum_{i=1}^{T} w_i(\hat{C}^{(j)})} \quad \text{with } x_i = \frac{\mathrm{z}_i}{ga\tau_i} \quad w_i = \frac{(ga\tau_i)^2}{g^2 a\tau_i \hat{C}^{(j)} + \sigma_R^2}. \tag{3.5}$$

It can be verified that in practice, if the weights are correctly initialized, the estimator remains almost unchanged after the first iteration. A correct initialization is $\hat{C}^{(0)} = \frac{\mathrm{z}_i}{ga\tau_i}$, i.e. the weights are computed directly from the input samples.

A first order approximation of the variance of the estimator after the first iteration can be computed as (c.f. Appendix C)

$$\mathrm{var}(\hat{C}^{(1)}) = \left[\sum_{i=1}^{T} \frac{(ga\tau_i)^2}{g^2 a\tau_i C + \sigma_R^2}\right]^{-1} + \mathrm{E}(o((\mathbf{z})^2)). \tag{3.6}$$

On the other hand, recall the CRLB formula introduced in (3.1)

$$\mathrm{CRLB} = \left[\sum_{i=1}^{T} \frac{(ga\tau_i)^2}{g^2 a\tau_i C + \sigma_R^2} + \frac{(g^2 a\tau_i)^2}{2(g^2 a\tau_i C + \sigma_R^2)^2}\right]^{-1}. \tag{3.7}$$

Hence the variance of $\hat{C}^{(1)}$ approaches the CRLB if

$$\frac{g^2 a^2 \tau_i^2}{g^2 a\tau_i C + \sigma_R^2} \gg \frac{g^4 a^2 \tau_i^2}{2(g^2 a\tau_i C + \sigma_R^2)^2} \quad \forall i = 1, \ldots, T, \tag{3.8}$$

which can be rewritten

$$a\tau_i C \gg \frac{1}{2} - \frac{\sigma_R^2}{g^2} \quad \forall i = 1, \ldots, T. \tag{3.9}$$

$a\tau_i C$ is the number of photons accumulated in the pixel, which is always larger than 1/2. This suggests that the iterative algorithm proposed by Granados et al. almost achieves the CRLB after the first iteration.

[2]Unlike Granados et al. we neglect the dark currents since their contribution to global noise is minimal for the exposures normally used in HDR image generation methods.

Figure 3.1: HDR image taken as ground-truth for the synthetic tests. Dynamic range 12.7 stops. From [18].

3.4. How close is the MLE from the CRLB: An experimental study. In §3.2 we showed that an efficient estimator for the irradiance does not exist under Model (2.2). Yet it is of interest to study how close to the CRLB existing methods perform. In particular, we are interested in the MLE since it was shown to outperform other methods [8, 11].

In the present section an experimental study is presented analyzing the performance of the MLE and five other methods widely known for the irradiance map estimation: Kirk and Andersen [8], Robertson et al. [9], Debevec and Malik [2], Mitsunaga and Nayar [16] and Reinhard et al. [4]. The experimental study is performed with synthetic data only since the knowledge of the exact ground-truth is imperative for the computation of the CRLB.

Synthetic data generation. Synthetic samples are generated from a HDR image taken as ground-truth assuming the pixel values follow Model (2.2). Figure 3.1 shows a tone mapped version of the ground-truth. Notice that a HDR image is taken as ground-truth in order to consider a realistic dynamic range, and to easily visualize the results. However, given that the estimation is done on a per-pixel basis, it only depends on each individual pixel irradiance, independently of the pixels' ordering or location in the image. In order to evaluate the dependence on the shutter speeds, 6 sets of exposure times are tested:

- Two sets of relatively long exposure times with 4 and 6 elements:
  - $\tau_{4L} = (1.3, 1/1.5, 1/3.1, 1/6.2)s$,
  - $\tau_{6L} = (1.3, 1/1.5, 1/3.1, 1/6.2, 1/12.4, 1/25)s$;
- Two sets of relatively short exposure times with 4 and 6 elements:
  - $\tau_{4S} = (1/50, 1/100, 1/200, 1/400)s$,
  - $\tau_{6S} = (1/50, 1/100, 1/200, 1/400, 1/600, 1/800)s$;
- Two sets of medium exposure times of 4 and 6 elements:
  - $\tau_{4M} = (1/12.4, 1/25, 1/50, 1/100)s$,
  - $\tau_{6M} = (1/6.2, 1/12.4, 1/25, 1/50, 1/100, 1/200)s$.

Two cameras are simulated[3]:

[3]The camera parameters were obtained using the calibration procedure by Granados et al. [11].

Camera A												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	**1.006**	0.049	26.377	91.765	1.094	0.090	4.054	5.414	4.967	6.023	6.662	9.207
τ_{6S}	**1.008**	0.051	27.233	93.224	1.120	0.113	11.77	19.32	17.41	24.59	25.75	41.15
τ_{4M}	**1.005**	0.057	32.111	106.16	1.104	0.090	2.650	3.368	3.507	4.440	4.420	6.640
τ_{6M}	**0.998**	0.065	56.609	137.49	1.150	0.096	2.063	2.286	6.068	12.12	4.891	12.28
τ_{4L}	**0.994**	0.082	68.893	164.33	1.146	0.075	1.116	0.256	1.030	0.075	1.045	0.067
τ_{6L}	**0.998**	0.069	87.462	181.00	1.175	0.087	1.173	0.252	1.147	0.346	1.054	0.077
Camera B												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	**1.013**	0.046	19.814	67.438	1.083	0.084	7.958	10.11	6.308	6.650	9.174	10.87
τ_{6S}	**1.013**	0.046	21.094	75.651	1.103	0.105	28.20	40.68	22.41	26.43	36.11	47.39
τ_{4M}	**1.011**	0.046	25.592	80.771	1.090	0.084	5.921	8.344	5.021	5.802	7.077	9.413
τ_{6M}	**1.008**	0.046	44.924	117.82	1.128	0.098	11.43	29.27	18.01	34.06	24.46	55.97
τ_{4L}	**1.003**	0.046	52.647	121.43	1.121	0.080	1.162	0.309	1.178	0.483	1.127	0.476
τ_{6L}	**1.004**	0.045	65.897	144.98	1.151	0.090	1.405	1.020	2.272	4.157	1.505	3.660

Table 3.1: Average and standard deviation of the ratio MSE/CRLB_{SAT} for all tested configurations. For the MLE the ratio is in the order of 1% above unity, with a standard deviation of about 5%, meaning that the MSE of the MLE is very close to the CRLB_{SAT}.

- Camera A, a Canon 7D set to ISO 200 ($g = 0.87; \sigma_R^2 = 31.6; \mu_R = 2046; z_{sat} = 14042$),
- Camera B, a Canon 400D set to ISO 400 ($g = 0.33; \sigma_R^2 = 6.2; \mu_R = 256; z_{sat} = 4056$).

The PRNU factors at each pixel are simulated following a Gaussian distribution with mean 1 and standard deviation 0.01 [11]. The dynamic range of the scene is 12.7 stops. The MLE estimation is computed using the implementation provided by Granados et al. [11]. Other irradiance estimation methods [9, 8, 2, 16, 4] are also tested in order to compare their performance against the CRLB. The estimations are repeated 1000 times for each irradiance level. The MSE and the variance of the estimators are computed from these 1000 repetitions.

Evaluation Procedure. Table 3.1 shows the average and standard deviation of the ratio between the MSE of each estimator and the CRLB for all the tested configurations. Since the Cramér-Rao bound is computed for unbiased estimators, the MSE equals the variance and the CRLB is a lower bound on the MSE. We opt to compute the ratio with respect to the MSE of the estimators and not to their variance since they are biased and thus they may have a low variance close to the CRLB with a large bias giving a large MSE. The main goal is to minimize the MSE and not only to decrease the variance.

The CRLB is computed for each pixel according to (3.1). In the CRLB computation we only take into account the exposure times producing non saturated samples, since in practice the samples corresponding to the other exposures would be saturated and therefore discarded (all methods discard the saturated samples). In the following, the CRLB computed from the non-saturated exposures only will be referred to as CRLB_{SAT}, while CRLB refers to the bound

Camera A												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	0.001	0.002	0.020	0.040	0.001	0.002	4.434	13.52	4.878	14.27	6.857	20.23
τ_{6S}	0.001	0.003	0.025	0.050	0.001	0.003	17.20	53.29	22.27	64.91	33.62	99.99
τ_{4M}	0.000	0.001	0.012	0.018	0.000	0.001	0.787	2.565	0.995	3.042	1.360	4.241
τ_{6M}	0.000	0.000	0.010	0.014	0.000	0.000	0.050	0.179	0.207	0.727	0.152	0.617
τ_{4L}	0.000	0.000	0.010	0.014	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
τ_{6L}	0.000	0.000	0.010	0.014	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.000
Camera B												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	0.005	0.040	0.093	0.383	0.005	0.039	71.02	379.3	65.00	343.0	89.29	470.1
τ_{6S}	0.004	0.023	0.128	0.584	0.004	0.023	322.4	1631	264.35	1279	400.9	1962
τ_{4M}	0.001	0.013	0.032	0.101	0.001	0.013	16.75	95.56	15.59	86.75	21.38	118.9
τ_{6M}	9e-05	0.000	0.011	0.017	9e-05	0.000	4.151	34.59	6.433	39.55	8.375	57.03
τ_{4L}	4e-06	2e-05	0.010	0.015	4e-06	2e-05	0.002	0.019	0.003	0.029	0.003	0.034
τ_{6L}	3e-06	1e-05	0.010	0.014	3e-06	1e-05	0.006	0.066	0.024	0.255	0.014	0.223

Table 3.2: Average and standard deviation of the ratio bias^2/C^2. Even if the MLE is known to be biased, its bias is negligible compared to the corresponding irradiance values.

computed considering all the exposures (ideal case without saturation).

3.4.1. Results.

Maximum Likelihood Estimation. For the MLE (Granados), in all cases the average ratio is in the order of 1% above unity, with a standard deviation of about 5%, meaning that the MSE of the MLE is very close to the CRLB_{SAT}. Moreover, it is of interest to remark that the bias of the MLE is negligible with respect to the corresponding irradiance values. Table 3.2 shows the expected ratio of the squared bias and the squared irradiance. Therefore, even if the MLE is biased, its bias is negligible and it can be fairly compared against the Cramér-Rao bound for unbiased estimators.

Figure 3.2 shows the comparison of the variance (left column) and the MSE (right column) curves against the CRLB (red) and the CRLB_{SAT} (black) curves as a function of the irradiance C for some of the tested configurations (similar results are found for all configurations). In agreement with the results presented on Tables 3.1 and 3.2, both the variance and the MSE of the MLE are very close to the CRLB_{SAT} (almost indistinguishable in these figures). It can be verified also that the bias is negligible for all irradiance values. The results are improved for the MLE when only considering high irradiance (c.f. Table 3.3). This is due to the fact that the irradiance range considered amply verifies condition (3.9).

Kirk and Andersen. The results presented on Table 3.1 for Kirk and Andersen's method show that the MSE for this method is quite larger than that obtained by the MLE. Actually,

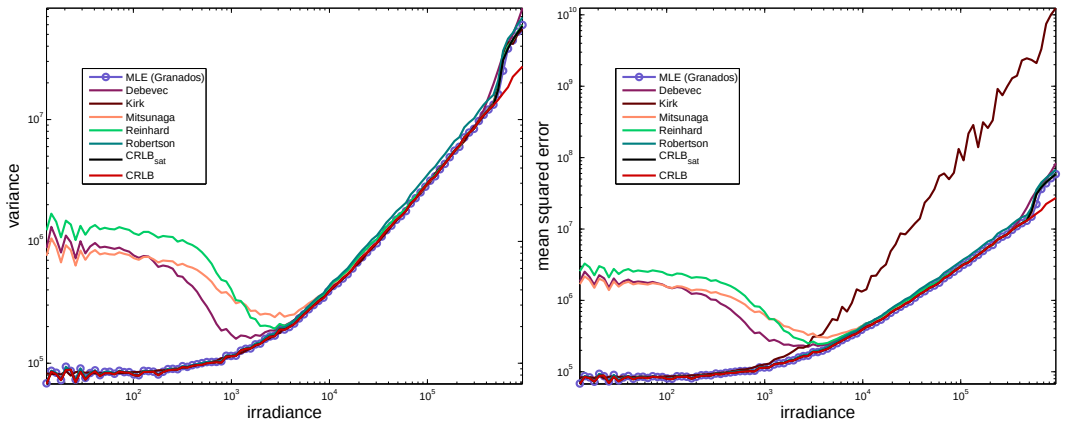

(a) Short exposure times

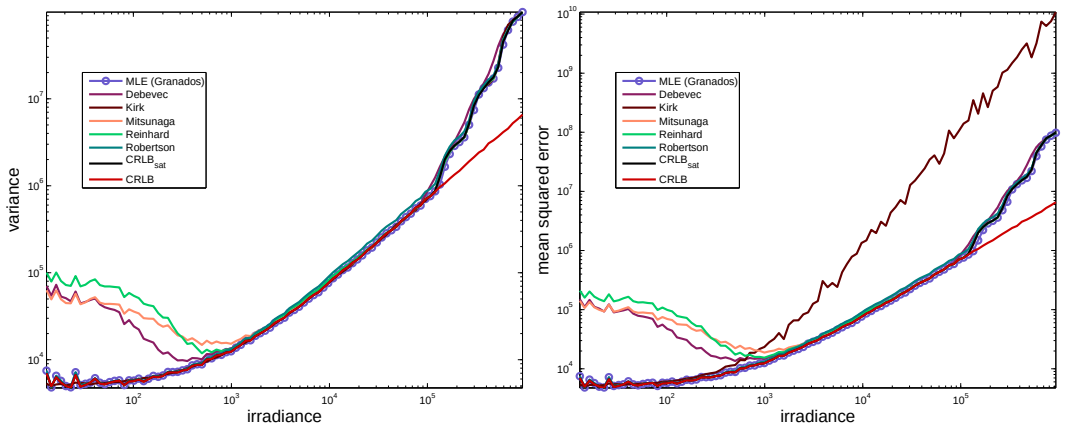

(b) Medium exposure times

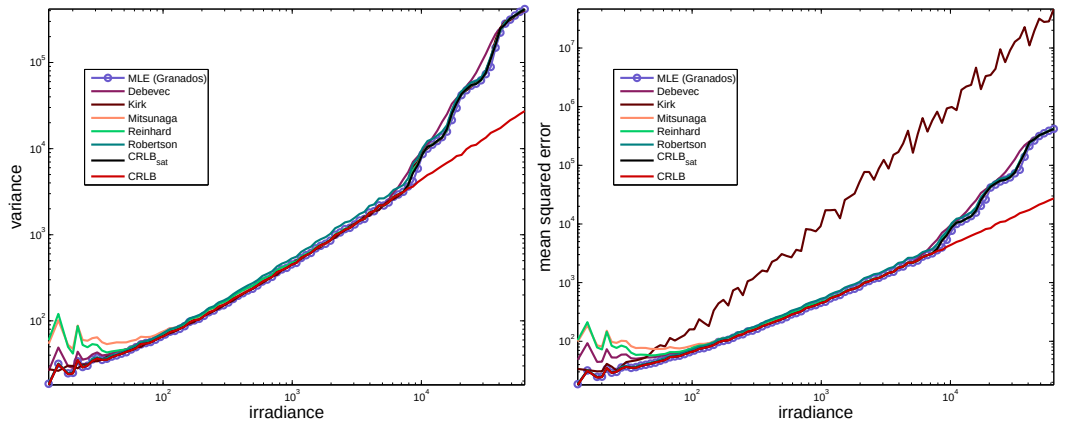

(c) Long exposure times

Figure 3.2: Comparison of the variance (left column) and MSE (right column) curves and the CRLB for Camera A with 4 exposure times. In all cases the variance of the MLE is very close to the CRLB.

Camera A												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	0.999	0.053	1.010	0.046	1.172	0.061	1.054	0.136	1.022	0.055	1.052	0.064
τ_{6S}	0.999	0.054	1.013	0.047	1.222	0.077	1.069	0.126	1.073	0.120	1.074	0.070
τ_{4M}	0.999	0.059	1.016	0.049	1.138	0.075	1.116	0.185	1.170	0.339	1.120	0.192
τ_{6M}	0.996	0.072	1.012	0.047	1.196	0.074	1.145	0.204	1.111	0.229	1.055	0.072
τ_{4L}	0.991	0.084	1.001	0.045	1.143	0.088	1.149	0.278	1.000	0.045	1.008	0.047
τ_{6L}	0.990	0.089	1.003	0.044	1.185	0.096	1.176	0.271	1.005	0.044	1.015	0.048
Camera B												
	Granados		Kirk		Robertson		Devebec		Mitsunaga		Reinhard	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
τ_{4S}	1.003	0.045	1.010	0.046	1.163	0.062	1.092	0.210	1.031	0.062	1.045	0.060
τ_{6S}	1.002	0.046	1.013	0.047	1.208	0.076	1.110	0.175	1.106	0.166	1.065	0.071
τ_{4M}	1.005	0.045	1.017	0.048	1.129	0.077	1.174	0.265	1.230	0.440	1.148	0.274
τ_{6M}	1.002	0.046	1.010	0.047	1.183	0.075	1.199	0.266	1.158	0.325	1.049	0.074
τ_{4L}	1.002	0.045	1.002	0.045	1.140	0.087	1.188	0.336	1.001	0.045	1.007	0.047
τ_{6L}	1.001	0.046	1.002	0.046	1.178	0.097	1.223	0.329	1.005	0.048	1.011	0.050

Table 3.3: Average and standard deviation of the ratio MSE/CRLB_{SAT} for high irradiance only. The simulation is performed without including the PRNU in Model (2.2). The results for all methods are quite close to the CRLB_{SAT}.

this large difference in performance is mainly introduced by the bias. The variance of the estimator is quite close to the CRLB_{SAT} as shown in Figure 3.2. The bias is caused by the PRNU, which is not taken into account in Kirk and Andersen's model.

Table 3.3 shows the results of the experiment without including the PRNU on the simulation of the data. Results highly improve in that case. Kirk and Andersen's estimator can be simply modified to include the PRNU and thus it can easily achieve the results presented in Table 3.3.

Results improve for Kirk and Andersen when only considering high irradiance since they approximate the variance directly from the input samples which are less noisy on high irradiance.

Robertson et al.. The results for the method by Robertson et al. are in the order of 10% above the CRLB_{SAT}. Its performance decreases for the longer exposure time sets. Their method takes as weights the squared exposure times. Hence, for the longer exposure time sets it gives low importance to exposures long enough to better contribute to the estimation.

These weighting schemes are highly sensitive to the noise in the input samples (e.g. for Mitsunaga and Nayar the weights are the pixel values) and therefore their performance is degraded in low irradiance where the SNR of the samples is lower.

Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.. The results obtained for Debevec and Malik, Mitsunaga and Nayar and Reinhard et al. are considerably far from the CRLB_{SAT}. This is mainly due to their poor performance at low irrradiances, as shown in Figure 3.2. Indeed, these weighting schemes are highly sensitive to noise in the input samples (e.g. for the

linear camera case, Mitsunaga and Nayar's weights are the pixel values) and therefore their performance is severely degraded for low irradiances where the SNR of the samples is lower.

On the other hand, as can be verified in Figure 3.2, their performance highly improves for the high irradiance range. Thus it is interesting to make a more local analysis of performance and compute the mean ratio without considering the low irradiance range. These results are presented in Table 3.3. The results are considerably better and in the order of those obtained by Kirk and Andersen. The results presented in Table 3.3 do not include the PRNU factor. However, this does not change the results for Debevec and Malik, Mitsunaga and Nayar and Reinhard et al. since their weights do not depend on the model parameters and the inverse of the camera response function is assumed to be known.

Debevec and Malik use a hat function to weigh the samples. For short exposures its weighting scheme is linearly increasing with the exposure. The decreasing part of the hat function is applied for longer exposures. Thus its performance decreases for the longer exposure times sets since a smaller weight is assigned to longer exposures which yet have the higher SNR. The opposite behavior is found for Mitsunaga and Nayar and Reinhard et al., since their weighting schemes are SNR based. Their performance increases for longer exposure time sets.

3.4.2. Summary of results. The experimental analysis carried out in this section shows that the MLE performs quite close to the CRLB_{SAT} for all irradiance values and for small sample sized datasets. It is important to remark that these results are for the small sample size case, since we have at most (none of the samples saturated) 4 or 6 samples for the estimation on each pixel. The performance of the MLE close to the CRLB was predicted in §3.3, where it was shown that the variance of the irradiance estimator almost attaint the bound after the first iteration.

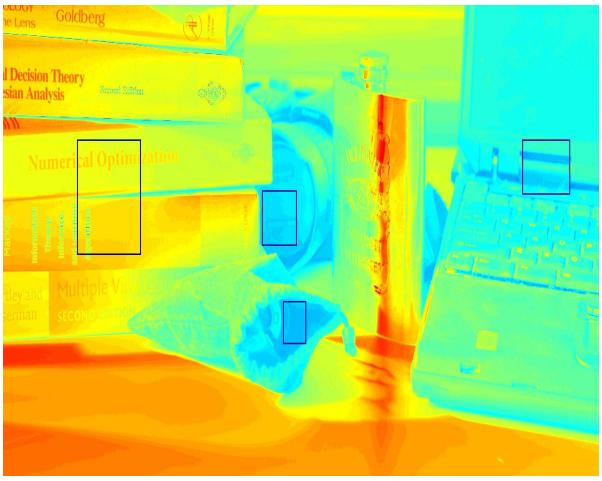

Figure 3.3: Logarithm of the ground-truth image. From [18].

Moreover, even if the $\hat{C}_{\text{MLE}}$ is not unbiased, its bias is negligible and its performance is very close to the best we can do among unbiased estimators. It would be interesting to evaluate

the performance of other biased estimators. However, there are no *a priori* hints on the parametric form that such a bias could take.This analysis also confirms the fact that the MLE outperforms other estimation methods, as observed in [8, 11].

Finally, it is important to remark that as expected, the CRLB using all samples is below the CRLB_{SAT}. This can be verified in Figure 3.2. Notice the difference between the two CRLB curves. This shows that discarding the saturated samples has a high impact on the bounds of estimation performance.

3.5. Visual quality of the estimators. In §3.4 we verified two facts: first the MLE performs quite close to the CRLB and second it outperforms other irradiance estimation methods. In this section we compare the visual quality of the results obtained by these methods. Is the performance difference among them noticeable in practice? To answer this question we compare their estimations in subregions of a test image using both synthetic and real data.

Synthetic Data. Figure 3.3 shows a subregion of the ground-truth image used for the synthetic experiments. The logarithm of the image is displayed for better visualization. The ground-truth image and the presented results correspond to the green channel of the image in Figure 3.1, yet the results are also valid for the red and blue channels. The results obtained for Camera A and τ_{6L} for the four marked subregions are displayed in Figures 3.4 to 3.6.

The first two results are for regions of low irradiance ($C < 10^2$). In agreement with the results presented in §3.4, the estimates produced by MLE, Kirk and Andersen and Robertson et al. are quite close and less noisy than those obtained by Debevec and Malik, Mitsunaga and Nayar and Reinhard et al. (see Figure 3.2c). In particular, the results of MLE and Robertson et al. are hardly distinguishable. The third subregion presented in Figure 3.5 corresponds to irradiance values in the order of 10^2. For that range, the estimations of MLE and Robertson et al. are still quite similar while the result for Kirk and Andersen degrades and is closer to that of Debevec and Malik, Mitsunaga and Nayar and Reinhard et al. The last subregion represents the behavior in high irradiance. The difference between all estimators, except Kirk and Andersen, is hardly noticeable. The result for Kirk and Andersen is quite degraded mostly because of the bias introduced by the PRNU. Yet, as previously stated, the PRNU can be easily incorporated to their estimator improving its performance.

Real Data. Tests using real data are carried out using the images provided by Granados et al. [10]. The set of images was acquired with a Canon PowerShot S5 camera, using exposures $\tau = (1/1.5, 1/6, 1/25, 1/100, 1/400, 1/1600)s$. The camera configuration is provided in [11]. The presented results correspond to the red channel, but similar results and considerations hold for the green and blue channels.

Figure 3.7 shows examples for irradiance values in the low - mid-level range ranges ($\sim 10^3$). As expected, the results for MLE, Kirk and Andersen and Robertson et al. are quite similar while the results for Debevec and Malik and Reinhard et al. are better than those of Mitsunaga and Nayar (see Figure 3.2b).

The examples in Figure 3.8 correspond to regions of mid-level and high irradiance ($\sim 10^{3.4}$ and $\sim 10^{4.4}$ respectively). In both cases all methods except Kirk and Andersen are expected to perform quite similarly (see Figure 3.2b). However, in practice quite similar results are found for all of them. The expected difference in performance for Kirk and Andersen is due to the PRNU. Thus the similarity of the results may be a consequence of a not sufficiently

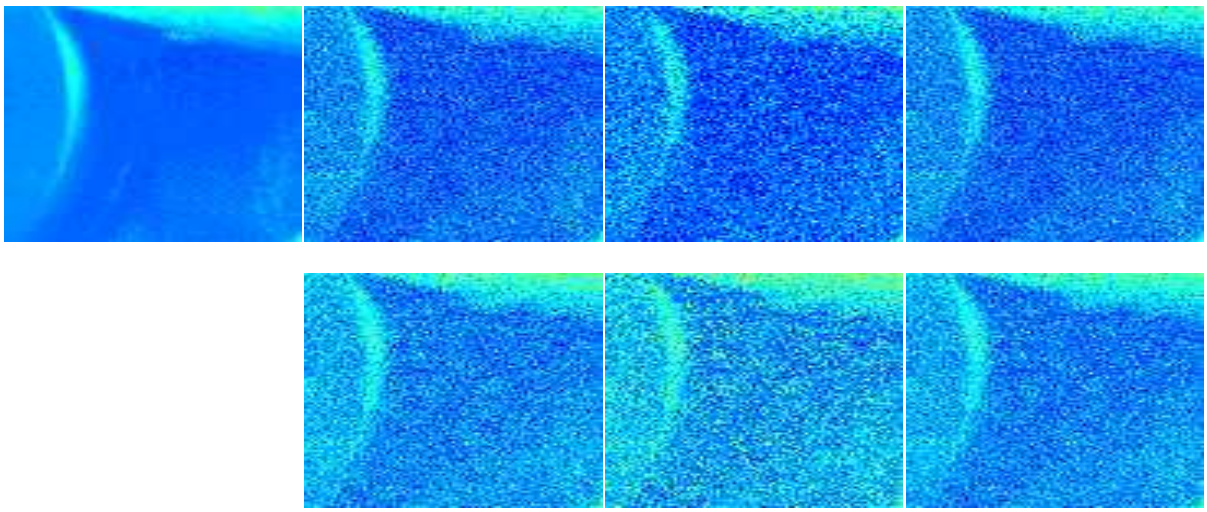

(a) **Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

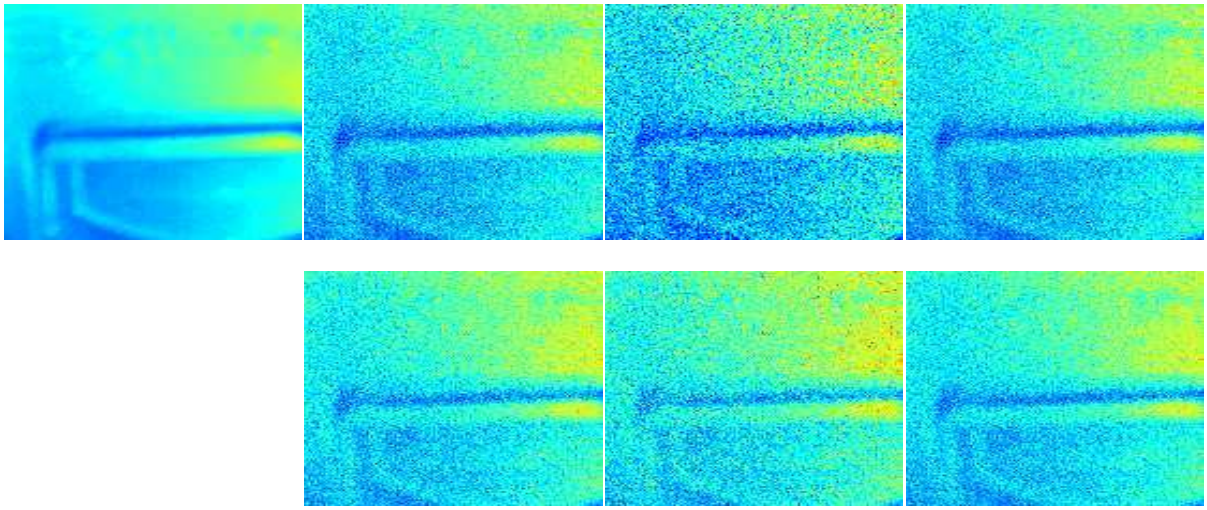

(b) **Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

Figure 3.4: **Synthetic data. Estimation in low irradiance range (~ 50).** The results for MLE, Kirk and Andersen and Robertson et al. are quite close and less noisy than those obtained by Debevec and Malik, Mitsunaga and Nayar and Reinhard et al. In particular, the results of MLE and Robertson et al. are hardly distinguishable.

precise estimation of the PRNU factors. In that case, the performance of all methods will be equally affected by the PRNU effect, as will be shown in §5.

4. Including saturation information. To the best of our knowledge, all approaches to irradiance estimation discard saturated samples. However, the experiments presented in §3.4

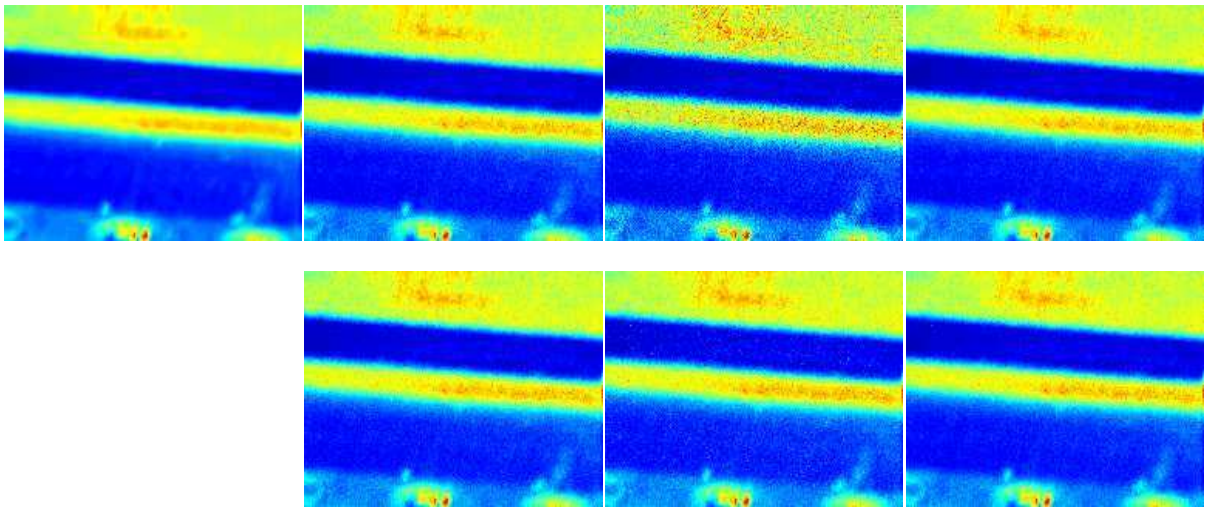

Figure 3.5: **Synthetic data. Estimation in mid-level irradiance range (~ 100). Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al. The results for MLE and Robertson et al. are still quite similar while the result for Kirk and Andersen degrades and is closer to that of Debevec and Malik, Mitsunaga and Nayar and Reinhard et al.

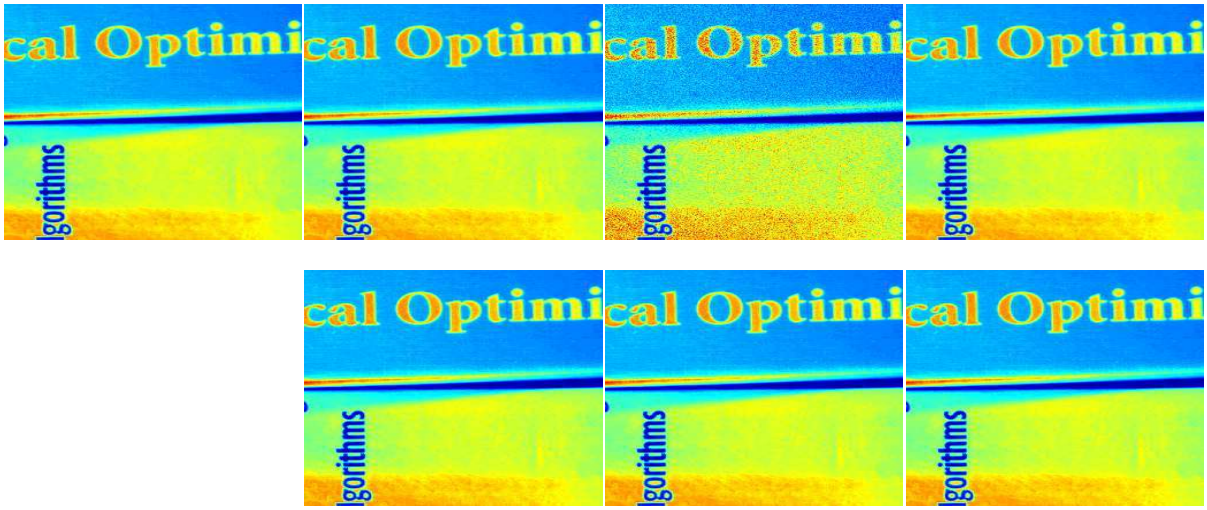

Figure 3.6: **Synthetic data. Estimation in high irradiance range ($\sim 10^{3.5}$).Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al. The difference between all estimators, except Kirk and Andersen, is hardly noticeable. The result for Kirk and Andersen is quite degraded mostly because of the bias introduced by the PRNU.

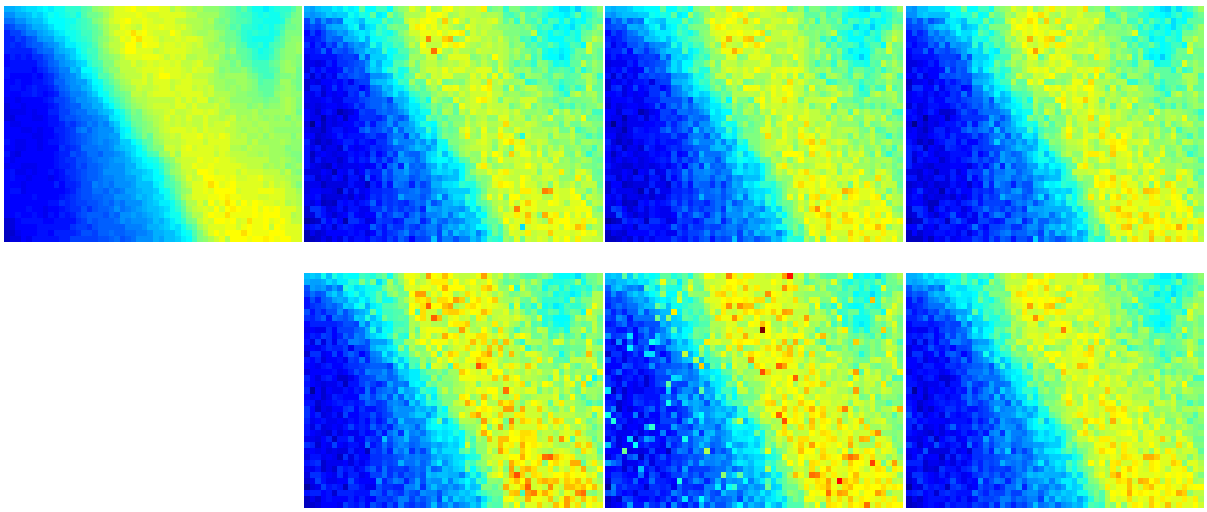

(a) **Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

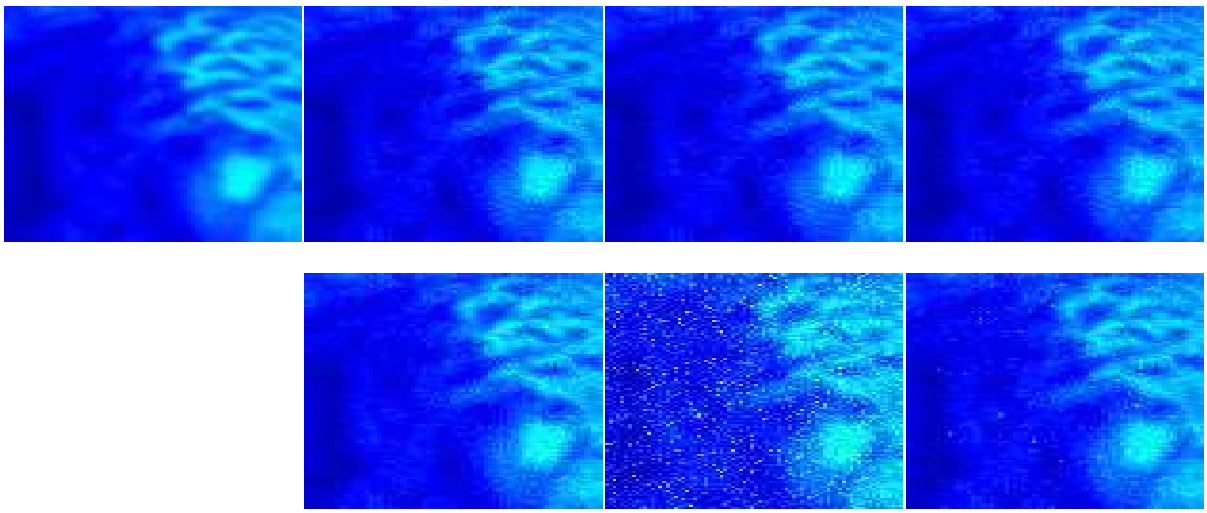

(b) **Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

Figure 3.7: **Real data. Estimation in low and mid-level irradiance range** ($\sim 10^3$) The results for MLE, Kirk and Andersen and Robertson et al. are quite similar while the results for Debevec and Malik and Reinhard et al. are better than that of Mitsunaga and Nayar.

show a great lost in performance when comparing the CRLB obtained from only non-saturated samples, with the one obtained from all the samples. We therefore propose a method to include information provided by the saturated samples in the irradiance estimation process.

4.1. Methodology. Although we do not know the exact value leading to a saturated sample, there is some useful information given by the fact that this value exceeds a given

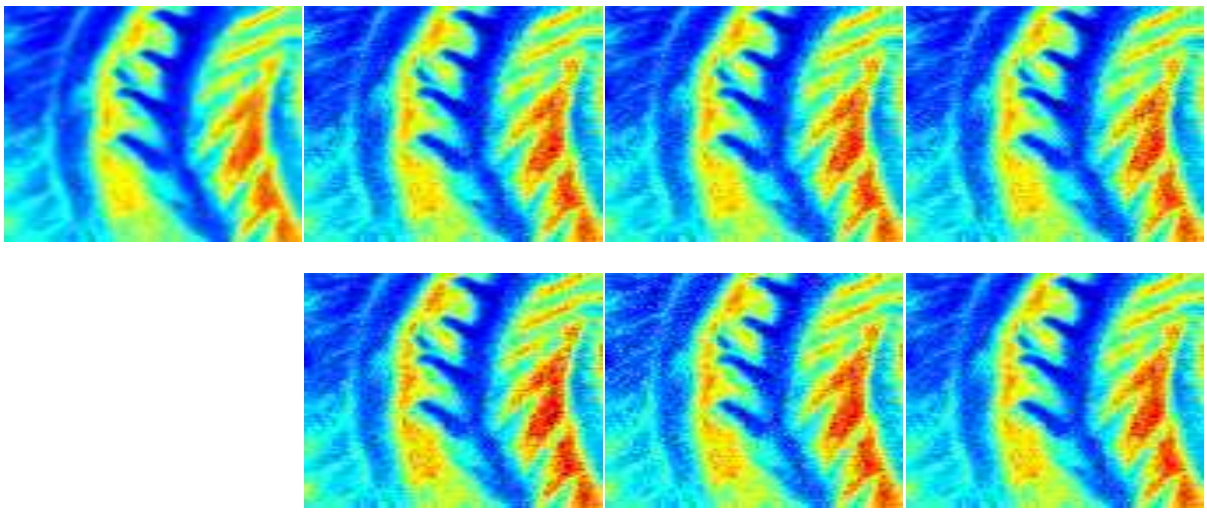

(a) **Estimation in mid-level irradiance range ($\sim 10^{3.4}$). Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

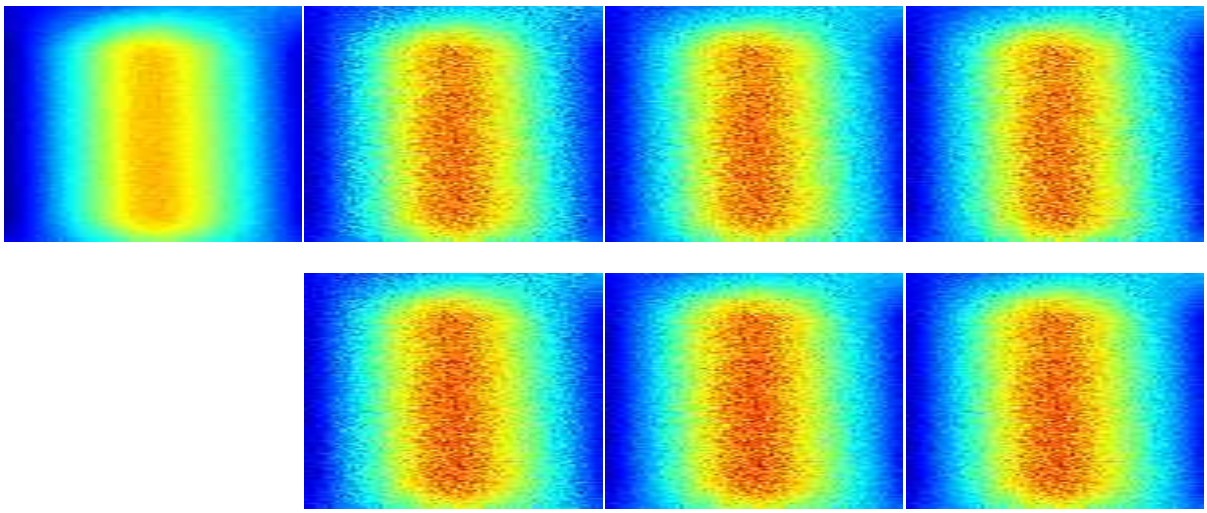

(b) **Estimation in high irradiance range ($\sim 10^{4.4}$). Top row:** Ground-truth, MLE, Kirk and Andersen, Robertson et al. **Bottom row:** Debevec and Malik, Mitsunaga and Nayar, Reinhard et al.

Figure 3.8: **Real data. Estimation in mid-level and high irradiance range.** In both cases (mid-level and high irradiance) all methods except Kirk and Andersen are expected to perform quite similarly. However, we find quite similar results for all of them. The similarity of the results may be a consequence of a not sufficiently precise estimation of the PRNU factors.

threshold z_{sat}. In this section, we develop a variation of the classical MLE that includes extra information provided by the saturated samples.

Let $\mathbf{Z}_1, \ldots, \mathbf{Z}_T$ be T independent but not identically distributed random variables, corresponding to the raw values observed at a given pixel p for the different exposure times

$\tau_1, \ldots, \tau_T$. Because of saturation, each $\mathbf{Z}_j$ can be written as $\mathbf{Z}_j = \min(\mathbf{X}_j, z_{sat})$, where $\mathbf{X}_j$ is a random variable following the law $N(\mu_j(C), \sigma_j^2(C))$, with $\mu_j(C) = ga\tau_j C$ and $\sigma_j^2(C) = g^2 a\tau_j C + \sigma_R^2$. Now, let us denote $p_j(\mathrm{z};\phi)$ the density of the law $N(\mu_j(\phi), \sigma_j^2(\phi))$. In the previous sections, saturated samples were discarded and an estimator of the irradiance C was obtained by maximizing the likelihood $\prod_{\mathrm{z}_j < z_{sat}} p_j(\mathrm{z}_j;\phi)$ as a function of ϕ. In order to take into account saturated samples, we propose to use instead full likelihood of the complete set of samples $\mathbf{z} = (\mathrm{z}_1, \ldots, \mathrm{z}_T)$, which writes

$$g(\mathbf{z}|\phi) = \prod_{j=1}^{T} p_j(\mathrm{z}_j;\phi)^{k_j} P_j(\phi)^{1-k_j}, \tag{4.1}$$

where $k_j = 1$ for non-saturated samples, and 0 otherwise. $P_j(\phi)$ is the probability of the j-th sample being saturated, given by

$$P_j(\phi) = \int_{z_{sat}}^{\infty} p_j(\mathrm{z};\phi) dz. \tag{4.2}$$

Observe that if non of the pixels saturate, $g(\mathbf{z}|\phi)$ is exactly the likelihood maximized by the MLE approach presented in the previous sections.

An efficient way to maximize the likelihood g is to use the EM algorithm [3]. In this setting, saturated samples are seen as censored data. Following Dempster et al. [3], we denote by $\mathbf{x} = (\mathrm{x}_1, \ldots, \mathrm{x}_T)$ the data that would have been observed if the camera could record beyond the saturation threshold. Each x_j is a realization of the random variable $\mathbf{X}_j$ defined above. We define the complete data likelihood of the problem as

$$h(\mathbf{x}|\phi) = \prod_{j=1}^{T} p_j(\mathrm{x}_j;\phi). \tag{4.3}$$

Now, observe that h cannot be maximized in practice because it relies on some censored data. The idea of the EM algorithm is to maximize instead the average value of $\log h$ knowing the observations $\mathbf{z}$, *i.e.* $\mathrm{E}_{\mathbf{x}}[\log h(\mathbf{x}|\phi)|\mathbf{z}, \phi]$. It can be shown that the value ϕ maximizing this expectation is exactly the same as the one maximizing $g(\mathbf{z}|\phi)$ (this is a classical result of the EM algorithm). The steps of the algorithm can finally be written

1. At iteration $p+1$ compute $Q(\phi|\phi^{(p)}) = \mathrm{E}_{\mathbf{x}}[\log h(\mathbf{x}|\phi)|\mathbf{z}, \phi^{(p)}]$;
2. Find $\phi^{(p+1)} = \max_\phi Q(\phi|\phi^{(p)})$.

The computation of the function $Q(\phi|\phi^{(p)})$ is provided in Appendix B.

4.2. The modified log-likelihood. The logarithm of the likelihood function $g(\mathbf{z}|\phi)$ is given by the sum of two terms,

$$\ln g(\mathbf{z}|\phi) = \underbrace{\sum_{j:k_j=1} \ln p_j(\mathrm{z}_j;\phi)}_{\text{term 1}} + \underbrace{\sum_{j:k_j=0} \ln P_j(\phi)}_{\text{term 2}}. \tag{4.4}$$

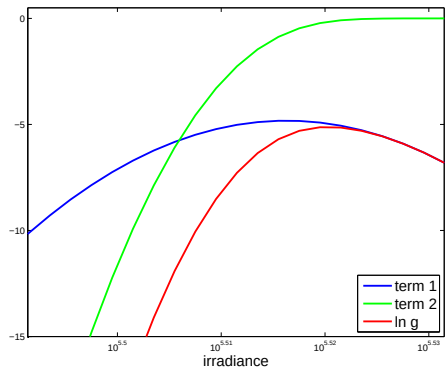

Figure 4.1: In blue the classical log-likelihood curve; in green function (4.5); in red the addition of the two curves: function $\ln g(\mathbf{y}|\phi)$.

The first term is the sum of the log-likelihood function evaluated at the non-saturated samples, i.e. the log-likelihood function of the MLE approach discarding the saturated samples (classical MLE approach). The second term is the sum of the logarithm of the probabilities of the saturated samples to be saturated. It can be written as

$$\sum_{j:k_j=0} \ln P_j(\phi) = \sum_{j:k_j=0} \ln \left[\int_{z_{sat}}^{\infty} \frac{1}{\sqrt{2\pi\sigma_j^2(\phi)}} \exp\left\{ -\frac{1}{2\sigma_j^2(\phi)} (\mathrm{z} - \mu_j(\phi))^2 \right\} d\mathrm{z} \right], \tag{4.5}$$

where $\mu_j(\phi) = ga\tau_j\phi$ and $\sigma_j^2(\phi) = g^2 a\tau_j\phi + \sigma_R^2$. The function $\sum_{j:k_j=0} \ln P_j(\phi)$ is increasing, takes negative values, and converges to zero for large ϕ, since large irradiance values have high probability to saturate. When we add the term (4.5) to the classical log-likelihood, its maximum will either: a) be almost unchanged, if it was reached on a value where (4.5) is quasi-zero; b) otherwise, be reached on a larger value.

Figure 4.1 shows an example of this modified log-likelihood $\ln g(\mathrm{z}|\phi)$ (in red) as a function of ϕ. In blue we find the classical log-likelihood curve; in green the additive function (4.5). In this example, the classical MLE is reached on a value ϕ where (4.5) is clearly negative. Therefore, when adding (4.5) and the classical log-likelihood, the maximum of the sum is reached for a larger ϕ. This example can be interpreted in the following way: if we use only the non-saturated samples we obtain the classical estimator $\hat{C}_{\mathrm{MLE}}$. However, knowing that we have saturated samples we increase $\hat{C}_{\mathrm{MLE}}$, since according to (4.5), the value $\hat{C}_{\mathrm{MLE}}$ is not sufficiently likely to saturate for the saturated exposure times. A higher value is representative of both, the non-saturated samples through $\hat{C}_{\mathrm{MLE}}$ and the saturated samples through the bias introduced by (4.5). Hence, the term (4.5) is adding a positive bias to $\hat{C}_{\mathrm{MLE}}$ when needed.

Irradiance values close to the minimum value that saturates for a given exposure may saturate due to noise. In those cases, incrementing the estimation with the saturation bias may degrade the result if the classical MLE was already accurate. Nevertheless, it is seen in practice that these cases rarely occur and in average adding the bias factor always improves the results.

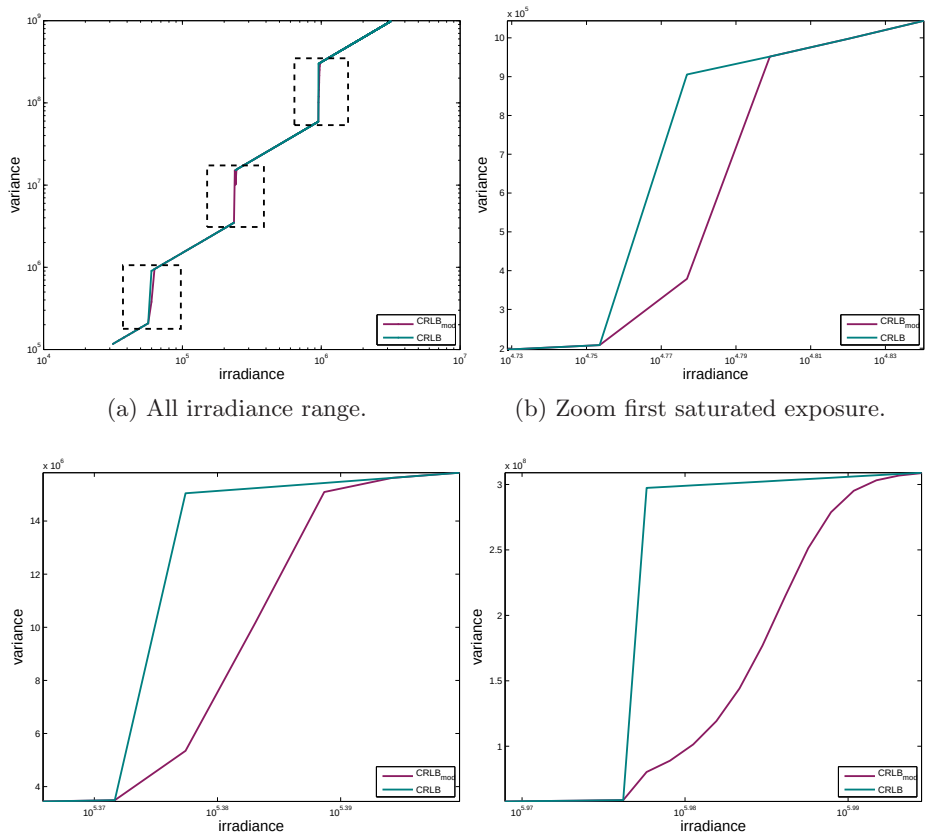

(a) All irradiance range. (b) Zoom first saturated exposure.

(c) Zoom second saturated exposure. (d) Zoom third saturated exposure.

Figure 4.2: Comparison of the CRLB obtained with the classical and the modified log-likelihoods. Both curves match except for the irradiance values in the transition between two saturation levels.

4.3. Modified Cramér-Rao Lower Bound. As performed in §3.2, the Cramér-Rao lower bound can be computed assuming the samples follow the modified log-likelihood function (4.4). The CRLB for the modified log-likelihood CRLB_{MOD} is given by

$$\text{CRLB}_{\text{MOD}} = \frac{-1}{\sum_{j:k_j=1} \mathrm{E}\left[\frac{\partial^2 \ln p(\mathrm{z}_j;C)}{\partial C^2}\right] + \sum_{j:k_j=0} \mathrm{E}\left[\frac{\partial^2 \ln P_j(C)}{\partial C^2}\right]}. \tag{4.6}$$

The first term in the denominator is the Fisher information for the classical log-likelihood. The second term can be computed as

$$\sum_{j:k_j=0} \mathrm{E}\left[\frac{\partial^2 \ln P_j(C)}{\partial C^2}\right] = \sum_{j:k_j=0} \frac{\int_{v_{sat}}^{\infty} \frac{\partial^2 A_j(C)}{\partial C^2} d\mathrm{z}}{\int_{v_{sat}}^{\infty} A_j(C) d\mathrm{z}} - \left[\frac{\int_{v_{sat}}^{\infty} \frac{\partial A_j(C)}{\partial C} d\mathrm{z}}{\int_{v_{sat}}^{\infty} A_j(C) d\mathrm{z}}\right]^2, \tag{4.7}$$

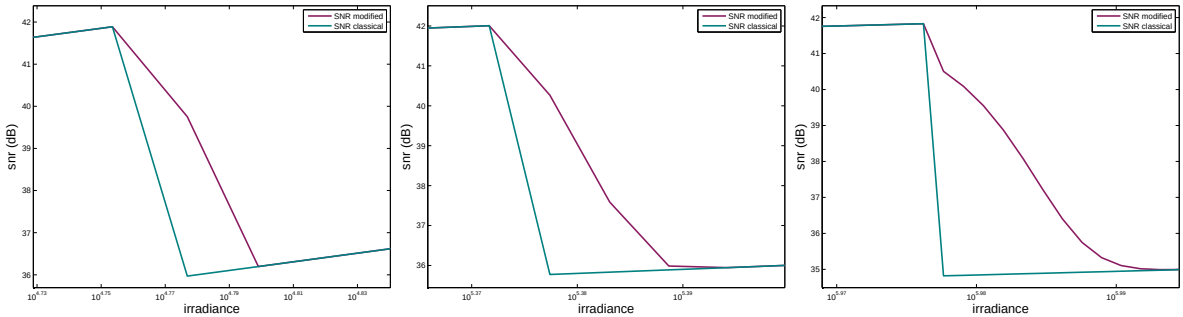

Figure 4.3: Comparison of the SNR obtained with the classical and the modified log-likelihoods. Both curves match except for the irradiance values in the transition between two saturation levels.

with

$$A_j(C) = \frac{1}{\sqrt{2\pi\sigma_j^2(C)}} \exp\left\{-\frac{1}{2\sigma_j^2(C)}(\mathrm{z} - \mu_j(C))^2\right\}. \tag{4.8}$$

Figure 4.2 shows a comparison of the CRLB for the estimators computed from the classical and modified log-likelihoods. Figure 4.2a shows the CRLB for all the irradiance range while figures 4.2b to 4.2d show the zoom of the transition between two saturation levels. We name saturation level n the irradiance range for which the longest n exposures produce saturated samples. Both CRLB curves match except in the transition zones between two saturation levels. This was expected since the contribution of the modified log-likelihood to the CRLB can be thought of as a measure of the information carried by the saturated samples. The higher information is in the region of higher uncertainty on whether pixels saturate or not. If we take the middle of the irradiance range between transition zones n and $n+1$, the corresponding pixels are highly likely to saturate for the longest n exposures and very unlikely to saturate for exposure $n+1$. Thus the information provided by the fact that the pixel saturates for the longest n exposures is not significant.
Figure 4.3 compares the SNR curves for the two cases. Once again the SNR differs only in the transition zones. Yet we observe that the estimation performance can be considerably increased for the concerned irradiance ranges.

4.4. Experiments.

Synthetic data generation. Synthetic samples are generated from a HDR image taken as ground-truth assuming the Model (2.2) for pixel values. The simulated camera is Camera A (c.f. §3.4). The exposure times are $\tau = (1/4.2, 1/16.8, 1/67.2, 1/268.8)s$.

Results. Figure 4.5 shows the results for four sub-regions of the ground-truth image, indicated on Figure 4.4. The pixels in these sub-regions have 2 or 3 saturated samples. The

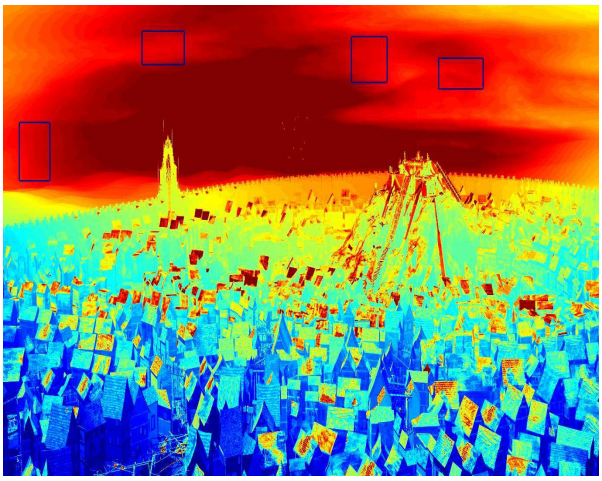

Figure 4.4: Logarithm of the ground-truth image. From [15] under Creative Commons License. The results for the four marked sub-regions are shown in Figure 4.5.

first column presents the ground-truth, the second the result obtained with the classical log-likelihood and the third the result obtained with the modified log-likelihood. In all cases, the information provided by the saturated samples improves the result. The improvement can also be verified in Table 5.1 where the values of PSNR for each sub-regions are presented.

5. Model parameters uncertainties and performance bounds. The results presented in §3.4 show that if the data follows Model (2.2), the MLE performs extremely well in estimating the irradiance. The estimation bias is negligible, and its variance gets very close to the CRLB_{SAT}. However, up to now we have assumed that the parameters that govern Model (2.2) are perfectly known. For practical purposes, this is of course not a realistic assumption: the gain factor, the readout noise mean and variance and the PRNU factors are unknown and have to be determined by means of a calibration procedure. Hence the model parameters are subject to uncertainties, whose impact in the irradiance estimation has to be quantified. The first part of this section is devoted to present an experimental study to assess how sensitive the MLE estimation is to variations in the model parameters.

In the second part of this section we concentrate on the consequences of ignoring the model parameter uncertainties when evaluating the performance of HDR generation techniques. Surprisingly, up to our knowledge, no previous work on HDR generation takes these uncertainties into account. Uncertainties in model parameters play a fundamental role in the creation of a ground-truth image from real, regular digital camera images. Hence, depending on how the ground-truth is generated, the impact of not considering model parameters uncertainties on the reported performance may vary. In any case, as we will see, ignoring these uncertainties may generally lead to overrated performances.

5.1. Sensitivity of the MLE to variations in the model parameters. The sensitivity analysis can only be performed with simulated images, since knowledge of the real parameter values is required. Synthetic data was generated according to Model (2.2), taking as ground-

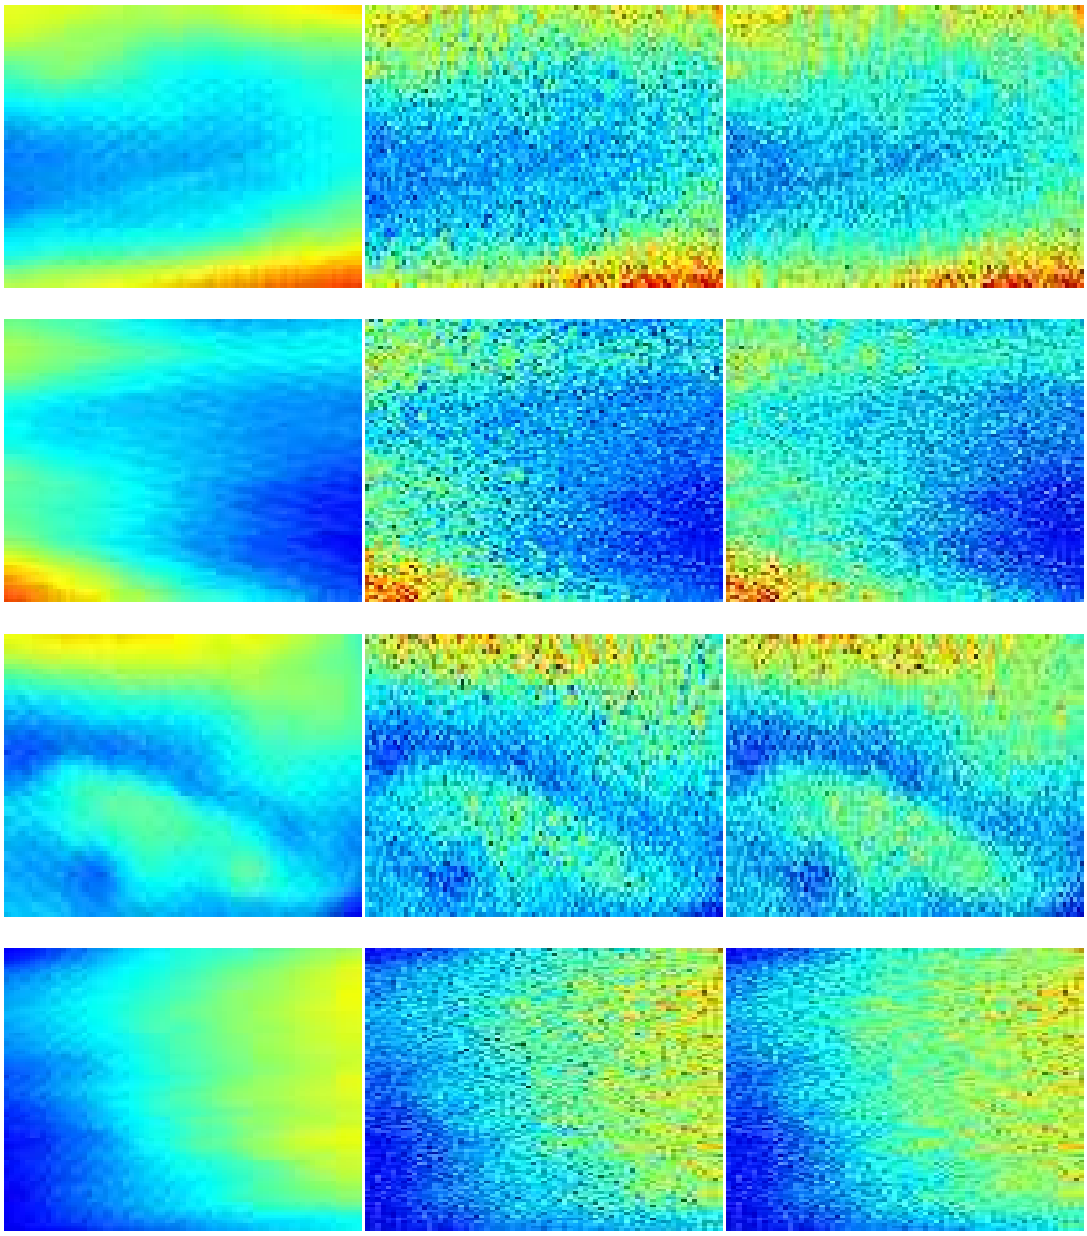

(j) Ground-truth (k) Classical log-likelihood (l) Modified log-likelihood

Figure 4.5: Results for four sub-regions of the ground-truth image in Figure 4.5. The pixels in these sub-regions have 2 or 3 saturated samples. **First column:** ground-truth. **Second column:** result obtained with the classical log-likelihood. **Third column:** result obtained with the modified log-likelihood. In all cases, the results are improved by the information provided by the saturated samples.

	PSNR (dB)			
	REGION 1	REGION 2	REGION 3	REGION 4
MODIFIED LOG-LIKELIHOOD	**37.7**	**38.7**	**39.3**	**38.1**
CLASSICAL LOG-LIKELIHOOD	36.9	37.6	38.3	37.0

Table 5.1: PSNR for the classical and modified likelihood for the four sub-regions marked in Figure 4.4.

truth the HDR image shown in Figure 3.1 and the set of exposure times τ_{M6}. The two sets of camera parameters presented in §3.4 (Cameras A and B) were tested. These parameters are considered the real model parameters and used to simulate the data. Then they are varied and the maximum likelihood estimation is performed using the *wrong* parameters. The ratio between the MSE and the CRLB_{SAT} is computed in order to compare the results with those obtained using the exact model parameters.

The first step in this study is to establish realistic ranges for the uncertainties in the model parameters, which is given by the variance of the estimators used to find each parameter. The model parameters are estimated at the camera calibration stage and the variance of the estimators clearly depends on the calibration procedure. To do so, we consider the calibration procedure proposed by Granados et al. [11]. They propose to compute the spatial mean and variance of a bias frame (a frame acquired with the cap on and with the shortest exposure) to find the mean and variance of the readout noise respectively. The gain and PRNU factors are computed using flat frames (frames acquired with uniform illumination). Using this calibration method we compute the variance of each model parameter.

Table 5.2 shows the standard deviation for the gain, the readout noise variance, the offset and the PRNU factors. The leftmost column represents the image size used to estimate the parameters. Several images are needed to estimate the PRNU factors, variable n represents the number images considered in each case. The real values of the parameters are those of camera A ($g = 0.87, \sigma_R^2 = 31.6, \mu_R = 2046$). The PRNU factors take Gaussian distributed values of mean 1 and variance 0.01. Similar results were obtained for Camera B.

The worst case standard deviations in Table 5.2 are used to define the variation range of the model parameters, namely: for g 1.5% of its value; 1.4% for σ_R^2; 2.7e-3% for μ_R and 1.3% for the PRNU factors. Because the variance of the estimator of the μ_R parameter is so small, the influence of the uncertainty on this parameter is not analyzed. Thus, the following variation ranges are considered for g, σ_R^2 and PRNU factors, respectively: $[0.985g, 1.015g]$ in steps of $0.001g$; $[0.99\sigma_R^2, 1.01\sigma_R^2]$, in steps of $0.001\sigma_R^2$; $[0.985a, 1.015a]$ in steps of $0.001a$

The influence of parameters uncertainties is analyzed individually: one of them is varied in its corresponding range, while the rest are kept fixed at their real values. For each tested (*wrong*) value, the estimation is repeated 1000 times for each irradiance level in order to compute the MSE.

Figure 5.1 shows the results for Camera A and Camera B configurations (left and right columns respectively). For each model parameter, the mean ratio of the MSE and the CRLB_{SAT} (blue) and the band of $\pm$ one standard deviation (dotted line) are shown as functions of the

Image size	g	σ_R^2	μ_R	a (n = 1)	a (n = 10)	a (n = 100)
100×100	0.013	0.447	0.056	0.013	0.004	0.001
200×200	0.006	0.224	0.028	0.013	0.004	0.001
500×500	0.003	0.089	0.011	0.013	0.004	0.001
1000×1000	0.001	0.045	0.006	0.013	0.004	0.001

Table 5.2: Standard deviation of the camera parameters estimated according to the procedure presented by Granados et al.[11]. Real values are those of Camera A ($g = 0.87, \sigma_R^2 = 31.6, \mu_R = 2046$).

parameter deviation. The unity level is shown for reference.

As expected the curves of mean ratio reach 1 when using the real parameter value (zero variation). A strong impact on the MLE is observed for both g and PRNU factor (this similar behavior is reasonable since they play a very similar role in the model). In both cases, for small deviations (1.5% of the real value) the MSE of the MLE almost doubles the CRLB_{SAT}. Regarding the readout noise variance σ_R^2, the result of the maximum likelihood estimation is not affected by variations in the considered range.

5.2. Model parameters uncertainties and performance evaluation. In order to evaluate the performance of a HDR image generation technique with real images, it is necessary to define a procedure to generate the ground-truth image. If the ground-truth is to be made from images taken with a regular digital camera (not cameras able to capture directly HDR images), a HDR image generation technique has to be chosen to generate the ground-truth image. It turns out that the precision of the model parameters estimates plays a key role in this ground-truth image generation process.

A simple way to generate the ground-truth is to define a set of exposure times, take several images for each exposure and take for each pixel the average value corresponding to the best non saturated exposure. This ground-truth is clearly unbiased and for a large enough number of images per exposure time, the noise should be considerably reduced. Instead of taking the best exposure only, Granados et al. [11] propose to combine the averaged images using the MLE with the weights computed from the variance of the images.

If the model parameters are correctly estimated, both ground-truths are quite accurate. Nonetheless, if the parameters are not accurate these ground-truths will be strongly biased. Moreover, it is not difficult to show that if the irradiance estimation is carried out using the same parameters as the ground-truth generation, the bias present in both computations will partially compensate and the result will seem better than it really is.

A synthetic experiment is performed to show this effect. Synthetic samples are generated as in the previous experiment, using Camera A configuration as real parameters and τ_{6L}. The gain spans the range $[0.985g, 1.015g]$ and the MLE is computed with the *wrong* parameters. Five different ground-truths are considered to compute the MSE: the real ground-truth (the image used to synthesize the samples), Granados's ground-truth obtained with the real parameters and with wrong parameters, the best exposure only ground-truth obtained with the real parameters and with wrong parameters. Figure 5.2 shows the ratio between the MSE of

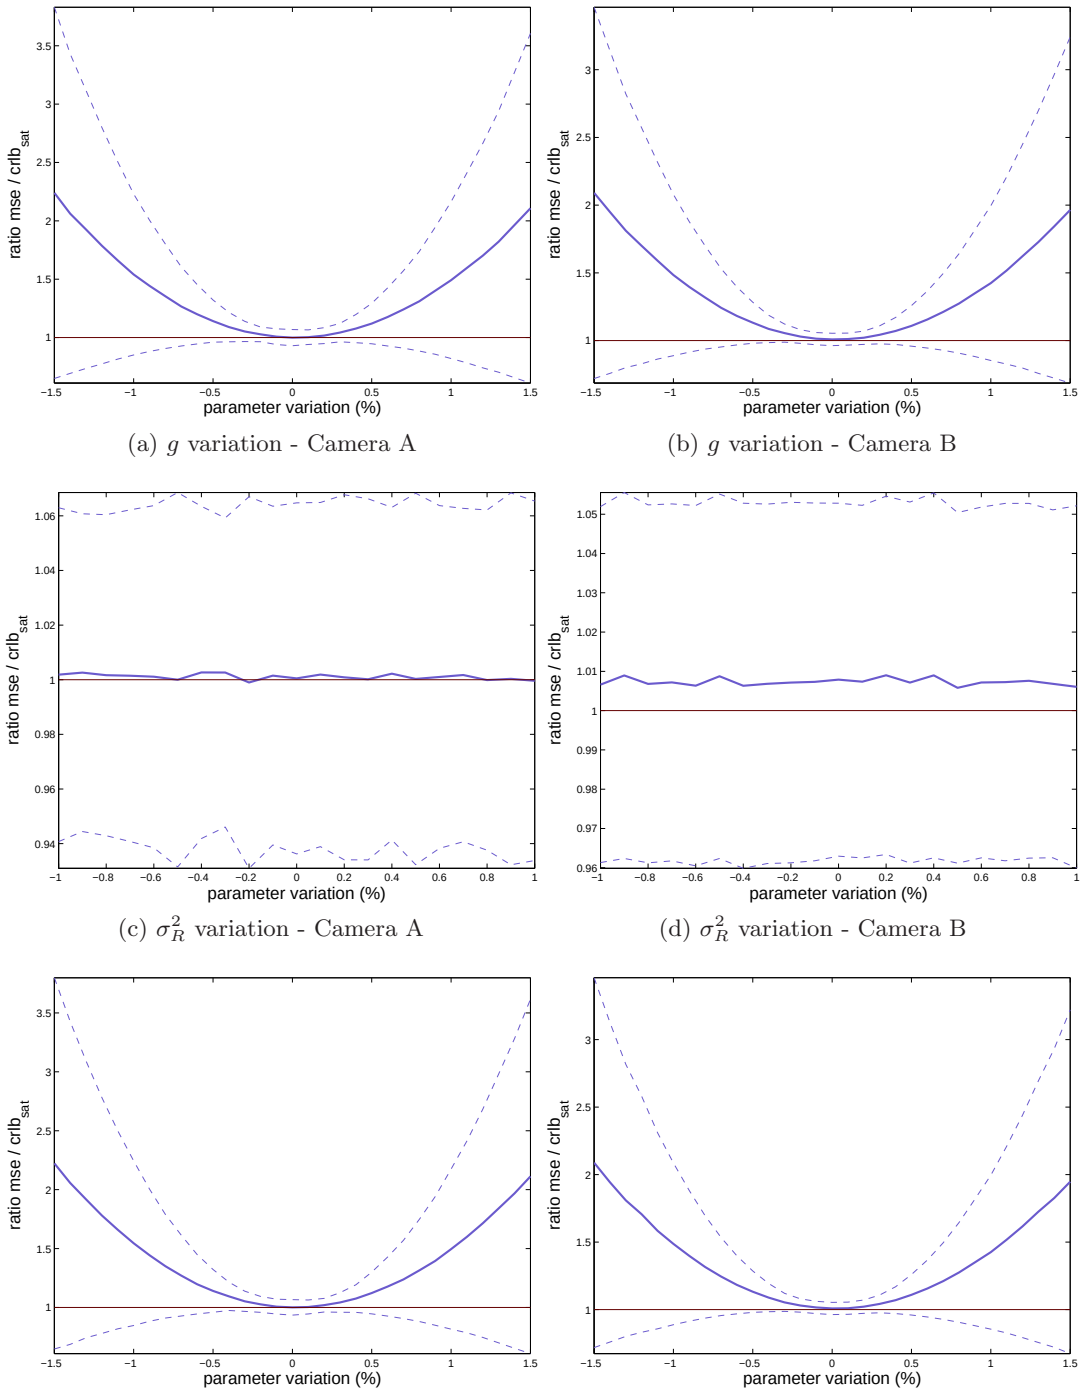

(a) g variation - Camera A

(b) g variation - Camera B

(c) σ_R^2 variation - Camera A

(d) σ_R^2 variation - Camera B

(e) PRNU variation - Camera A

(f) PRNU variation - Camera B

Figure 5.1: Dependence of the performance of the MLE estimator with the uncertainty in model parameters. **From top to bottom:** dependence on g, σ_R^2 and a. A great dependence is found with both g and a. On the contrary, the variation for σ_R^2 is negligible.

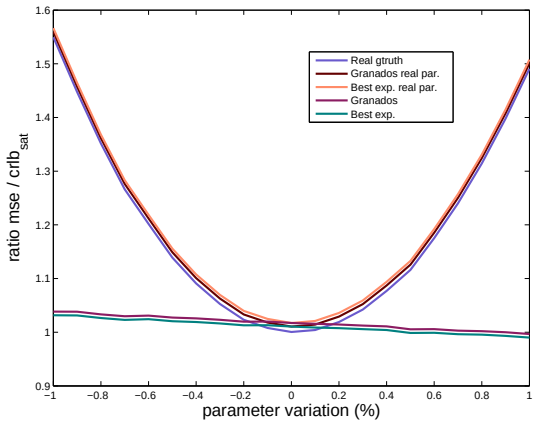

Figure 5.2: Dependence of the ground-truth computation with the uncertainty in model parameters. For small variations of the parameter the results of all computed ground-truths are almost equivalent. Yet for variations above 0.2% the performance announced by both ground-truths computed with the *wrong* parameters is highly superior than the real performance.

the MLE and the CRLB_{SAT} as a function of the variation in the gain. Each curve shows the results for a different ground-truth. For small variations of g the result given for the computed ground-truths is quite accurate. The ratio for the real ground-truth is below the other ratios, meaning that the result is actually better than claimed by the computed ground-truths. On the contrary, for deviations of g above 0.2%, the performance announced by both ground-truths computed with the *wrong* parameters is highly superior than the real performance.

6. Conclusions. In this paper we have presented a study of the performance bounds of the HDR estimation problem, and we have analyzed the performance of current state-of-the-art estimation methods. This study shows that, to a first order approximation, the MLE is efficient even when using a very reduced number of samples. While its value can be computed numerically by iterative procedures, an approximation of the MLE cannot be directly derived since no closed-form exists. However, we observe that replacing the variance of each sample by its empirical value, yields to a closed-form which is extremely close to the MLE. Then we show that the first order Taylor expansion of the variance of this closed-form estimator, exhibits a negligible difference with Cramér-Rao bound, for all irradiance values and for any number of samples. This result explains why, as previously claimed based on experimental evidence [11], the MLE outperforms other estimation methods. An interesting problem would be to extend this performance analysis to the choice of optimal exposure times.

As a second contribution, we have proposed a method that differs from all methods in the literature, in the sense that it integrates the information provided by saturated samples in the estimation process. The proposed approach follows closely the EM algorithm, in the version that considers censored data. Results confirm that saturated samples carry useful information for irradiance estimation, that allows to improve the irradiance estimation near the saturation values, without degrading the estimation in other irradiance ranges. Improvements in the

order of 1 dB are found in experimental examples.

Finally, we have raised a delicate point that had not been addressed in previous studies on HDR estimation. We have shown that small errors in the calibration of camera parameters may severely degrade the estimation. In particular, when working with real data, a very accurate camera calibration is needed in order to obtain a reliable ground-truth. Otherwise, results may appear much better than they are in reality. A natural continuation of this work is to derive irradiance estimators that are more robust to uncertainties on the camera parameters.

Acknowledgments. The authors would like to thank Miguel Granados for kindly providing them with the code and for helpful contributions.

Appendix A. CRLB computation. In this section we compute the Cramér-Rao lower bound (CRLB) for the estimation of the irradiance parameter C from T independent variables $\mathrm{z}_1, \ldots, \mathrm{z}_T$ corresponding to pixel values for exposure times $\tau_1, \ldots, \tau_T$. We assume the z_i variables are normally distributed with $\mu_i = ga\tau_i C$ and $\sigma_i^2 = g^2 a\tau_i C + \sigma_R^2$. The log-likelihood function is

$$\ln p(\mathbf{z}, C) = -\frac{1}{2}\sum_{i=1}^{T} \ln(g^2 a\tau_i C + \sigma_R^2) + \frac{(\mathrm{z}_i - ga\tau_i C)^2}{g^2 a\tau_i C + \sigma_R^2}. \tag{A.1}$$

The second derivative is given by

$$\frac{\partial^2 \ln p(\mathbf{z}, C)}{\partial C^2} = \frac{1}{2}\sum_{i=1}^{T} \frac{(ga\tau_i)^2 (g^4 a\tau_i C - 2g^2 \mathrm{z}_i^2 - 4g\sigma_R^2 \mathrm{z}_i + g^2\sigma_R^2 - 2)\sigma_R^4)}{(g^2 a\tau_i C + \sigma_R^2)^3}. \tag{A.2}$$

The CRLB is defined as

$$\mathrm{CRLB} = \frac{1}{I(C)} = \frac{1}{-E\left[\frac{\partial^2 \ln p(\mathbf{z}, C)}{\partial C^2}\right]} \tag{A.3}$$

Using the linearity of the expectation, and that $\mathrm{E}(\mathrm{z}_i) = ga\tau_i C$, $\mathrm{var}(\mathrm{z}_i) = g^2 a\tau_i C + \sigma_R^2$, it follows that

$$E\left[\frac{\partial^2 \ln p(\mathbf{z}, C)}{\partial C^2}\right] = -\left(\frac{g^2}{2}\sum_{i=1}^{T} \frac{2a^2\tau_i^2}{g^2 a\tau_i C + \sigma_R^2} + \frac{(ga\tau_i)^2}{(g^2 a\tau_i C + \sigma_R^2)^2}\right). \tag{A.4}$$

Then,

$$\mathrm{CRLB} = \left[\sum_{i=1}^{T} \frac{(ga\tau_i)^2}{g^2 a\tau_i C + \sigma_R^2} + \frac{(g^2 a\tau_i)^2}{2(g^2 a\tau_i C + \sigma_R^2)^2}\right]^{-1} \tag{A.5}$$

Appendix B. $Q(\phi|\phi^{(p)})$ computation. To compute $Q(\phi|\phi^{(p)})$ we separate the product (4.3) onto two terms, one including the known samples (NS) and the other considering the saturated samples (S).

$$\begin{aligned} Q(\phi|\phi^{(p)}) &= \mathrm{E}_{\mathbf{x}}[\log h(\mathbf{x}|\phi)|\mathbf{z}, \phi^{(p)}] \\ &\overset{(a)}{=} \sum_{j\in NS} \log p(\mathrm{z}_j|\phi_j) + \sum_{j\in S} \mathrm{E}_{\mathrm{z}_j}\left[\log p(\mathrm{z}_j|\phi_j)|\mathbf{z}, \phi^{(p)}\right], \end{aligned} \tag{B.1}$$

(B.2)

(a) Since z_j is known for $j \in NS$, the value $\log p(z_j|\phi_j)$ is no longer a random variable but a deterministic value.
For $j \in S$ we do not know z_j exact value but know its distribution

$$\mathrm{E}_{z_j}\left[\log p(z_j|\phi_j)|\mathbf{z},\phi^{(p)}\right] = \int_{z_{sat}}^{\infty} \log p(z_j|\phi_j)\left(\frac{p(z_j|\phi_j^{(p)})}{P_j}\right) dz_j \tag{B.3}$$

$$= -\frac{1}{2}\left[\ln(2\pi\sigma_j^2) + \frac{\mu_j^2}{\sigma_j^2} + \frac{1}{\sigma_j^2}\int_{z_{sat}}^{\infty} z_j^2\left(\frac{p(z_j|\phi_j^{(p)})}{P_j}\right) dz_j \ldots \right.$$
$$\left. -\frac{2\mu_j}{\sigma_j^2}\int_{z_{sat}}^{\infty} z_j\left(\frac{p(z_j|\phi_j^{(p)})}{P_j}\right) dz_j\right]. \tag{B.4}$$

We now compute

$$\int_{z_{sat}}^{\infty} z_j^2\left(\frac{p(z_j|\phi_j^{(p)})}{P_j}\right) dz_j \overset{(a)}{=} \mathrm{E}[z_j^2] \tag{B.5}$$

$$= \sigma_j^{2(p)}[1-\delta(\alpha)] + [\mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha)]^2, \tag{B.6}$$

$$\tag{B.7}$$

(a) z following the truncated normal distribution $N(\mu_j^{(p)}, \sigma_j^{2(p)})$ in the interval $[z_{sat}, \infty)$. Hence, with

$$\alpha = \frac{z_{sat} - \mu_j^{(p)}}{\sigma_j^{(p)}} \quad \lambda(\alpha) = \frac{\phi(\alpha)}{1-\Phi(\alpha)} \quad \delta(\alpha) = \lambda(\alpha)(\lambda(\alpha) - \alpha), \tag{B.8}$$

$\phi(\cdot)$ is the probability density function of the standard normal distribution and $\Phi(\cdot)$ is its cumulative distribution function .
We continue with

$$\int_{z_{sat}}^{\infty} z_j\left(\frac{p(z_j|\phi_j^{(p)})}{P_j}\right) dz_j \overset{(b)}{=} \mathrm{E}[z_j] \tag{B.9}$$

$$= \mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha), \tag{B.10}$$

(b) z following the truncated normal distribution $N(\mu_j^{(p)}, \sigma_j^{2(p)})$ in the interval $[z_{sat}, \infty)$. Hence,

$$\mathrm{E}_{z_j}\left[\log p(z_j|\phi_j)|\mathbf{z},\phi^{(p)}\right] = -\frac{1}{2}\left[\ln(2\pi\sigma_j^2) + \frac{1}{\sigma_j^2}\left\{\mu_j^2 - 2(\mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha))\mu_j + \ldots \right.\right.$$
$$\left.\left. +\sigma_j^{2(p)}(1-\delta(\alpha)) + (\mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha))^2\right\}\right]. \tag{B.11}$$

Thus, from (B.2) and (B.11)

$$Q(\phi|\phi^{(p)}) = \sum_{j\in NS} \log p(\mathrm{z}_j|\phi_j) - \frac{1}{2}\sum_{j\in S}\left[\ln(2\pi\sigma_j^2) + \frac{1}{\sigma_j^2}\left\{\mu_j^2 - 2(\mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha))\mu_j + \ldots\right.\right. \tag{B.12}$$

$$\left.\left. +\sigma_j^{2(p)}(1-\delta(\alpha)) + (\mu_j^{(p)} + \sigma_j^{(p)}\lambda(\alpha))^2\right\}\right]. \tag{B.13}$$

Appendix C. MLE variance computation. Let $\mathbf{y} = (y_1, \ldots, y_T)$ be a random vector, with $y_i = \mathrm{z}_i$. Thus $\mathbf{y} \sim N(\boldsymbol{\mu}, \Sigma)$ with

$$\boldsymbol{\mu} = (\mu_1, \ldots, \mu_T), \quad \mu_i = g\tau_i aC \quad \forall\, i = 1, \ldots, T, \tag{C.1}$$

$$\Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_T^2), \quad \sigma_i^2 = g^2\tau_i aC + \sigma_R^2 \quad \forall\, i = 1, \ldots, T. \tag{C.2}$$

Consider the function

$$h(\mathbf{y}) = \frac{\sum_{i=1}^T \frac{g\tau_i a y_i}{g y_i + \sigma_R^2}}{\sum_{i=1}^T \frac{(g\tau_i a)^2}{g y_i + \sigma_R^2}}. \tag{C.3}$$

Being $\mathbf{y}$ Gaussian distributed, and noting that h is continously differentiable and that

$$\sum_{j=1}^T \sum_{k=1}^T \sigma_{jk} \frac{\partial h(\boldsymbol{\mu})}{\partial y_j}\frac{\partial h(\boldsymbol{\mu})}{\partial y_k} > 0, \tag{C.4}$$

it follows from the Delta method theorem [17] that the variance of $h(\mathbf{y})$ can be computed to a first order approximation as

$$\mathrm{var}(h(\mathbf{y})) = \sum_{i=1}^T \left[\left.\frac{\partial h(\mathbf{y})}{\partial y_i}\right|_{\mathbf{y}=\mu_{\mathbf{y}}}\right]^2 \sigma_{y_i}^2 + \mathrm{E}(o((\mathbf{y}-\boldsymbol{\mu})^2)) \tag{C.5}$$

$$= \sum_{i=1}^T \left[\frac{g\tau_i a}{g^2\tau_i aC + \sigma_R^2}\right]^2 \left[\frac{(g\tau_i a)^2}{g^2\tau_i aC + \sigma_R^2}\right]^{-2} (g^2\tau_i aC + \sigma_R^2) + \mathrm{E}(o((\mathbf{y}-\boldsymbol{\mu})^2)) \tag{C.6}$$

$$\approx \left[\frac{(g\tau_i a)^2}{g^2\tau_i aC + \sigma_R^2}\right]^{-1}. \tag{C.7}$$

REFERENCES

[1] Y. Gousseau P. Muse C. Aguerrebere, J. Delon, *Study of the digital camera acquisition process and statistical modeling of the sensor raw data*, Preprint HAL http://hal.archives-ouvertes.fr/docs/00/73/35/38/PDF/camera_model.pdf, (2012).

[2] P. E. Debevec and J. Malik, *Recovering high dynamic range radiance maps from photographs*, in SIGGRAPH, 1997, pp. 369–378.

[3] A. P. Dempster, N. M. Laird, and D. B. Rubin, *Maximum likelihood from incomplete data via the EM algorithm*, Journal of the Royal Statistical Society, Series B, 39 (1977), pp. 1–38.

[4] S. N. Pattanaik E. Reinhard, G. Ward and P. E. Debevec, *High Dynamic Range Imaging - Acquisition, Display, and Image-Based Lighting*, Morgan Kaufmann, 2005.

[5] G.E. Healey and R. Kondepudy, *Radiometric CCD camera calibration and noise estimation*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 16 (1994), pp. 267 –276.

[6] A. Nemirovsky I. Brouk and Y. Nemirovsky, *Analysis of noise in CMOS image sensor*, in Microwaves, Communications, Antennas and Electronic Systems, 2008. COMCAS 2008. IEEE International Conference on, may 2008, pp. 1 –8.

[7] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

[8] K. Kirk and H. J. Andersen, *Noise characterization of weighting schemes for combination of multiple exposures*, in BMVC, 2006, pp. 1129–1138.

[9] S. Borman M. A. Robertson and R. L. Stevenson, *Estimation-theoretic approach to dynamic range enhancement using multiple exposures*, J. Electronic Imaging, 12 (2003), pp. 219–228.

[10] M. Wand C. Theobalt H. P. Seidel M. Granados, B. Ajdin and H. P. A. Lensch, *Optimal HDR reconstruction with linear digital cameras*. http://www.mpi-inf.mpg.de/~granados/projects/opthdr/index.html. Accessed: 15/08/2012.

[11] ———, *Optimal HDR reconstruction with linear digital cameras*, in CVPR, 2010, pp. 215–222.

[12] R. Mancini, *Op Amps for everyone*, Texas Instruments, Aug 2002.

[13] S. Mann and R. W. Picard, *On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures*, in Proceedings of IS&T, 1995, pp. 442–448.

[14] E. Martinec, *Noise, dynamic range and bit depth in digital slrs. thermal noise*. http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/#shotnoise. Accessed: 03/08/2012.

[15] Xiph.org Test Media, *Xiph.org test media*. http://media.xiph.org/sintel/sintel-4k-tiff16/. Accessed: 15/08/2012.

[16] T. Mitsunaga and S. K. Nayar, *Radiometric self calibration*, in CVPR, 1999, pp. 1374–1380.

[17] J. A. Rice, *Mathematical Statistics and Data Analysis*, Duxbury Press, 1995.

[18] F. Durand S. Hasinoff and W. Freeman, *Noise-optimal capture for high dynamic range photography*. http://people.csail.mit.edu/hasinoff/hdrnoise/. Accessed: 13/08/2012.

[19] F. Durand S. W. Hasinoff and W. T. Freeman, *Noise-optimal capture for high dynamic range photography*, in CVPR, 2010, pp. 553–560.

[20] A. J. P. Theuwissen, *Solid-State Imaging with Charge-Coupled Devices*, (1996), pp. 94–108.

[21] V. Ramesh Y. Tsin and T. Kanade, *Statistical calibration of the CCD imaging process*, in ICCV, 2001, pp. 480–487.