

RESEARCH

Open Access

Realistic camera noise modeling with application to improved HDR synthesis

Bart Goossens*, Hiệp Luong, Jan Aelterman, Aleksandra Pižurica and Wilfried Philips

Abstract

Due to the ongoing miniaturization of digital camera sensors and the steady increase of the “number of megapixels”, individual sensor elements of the camera become more sensitive to noise, even deteriorating the final image quality. To go around this problem, sophisticated processing algorithms in the devices, can help to maximally exploit the knowledge on the sensor characteristics (e.g., in terms of noise), and offer a better image reconstruction. Although a lot of research focuses on rather simplistic noise models, such as stationary additive white Gaussian noise, only limited attention has gone to more realistic digital camera noise models. In this article, we first present a digital camera noise model that takes several processing steps in the camera into account, such as sensor signal amplification, clipping, post-processing, ... We then apply this noise model to the reconstruction problem of high dynamic range (HDR) images from a small set of low dynamic range (LDR) exposures of a static scene. In literature, HDR reconstruction is mostly performed by computing a weighted average, in which the weights are directly related to the observer pixel intensities of the LDR image. In this work, we derive a Bayesian probabilistic formulation of a weighting function that is near-optimal in the MSE sense (or SNR sense) of the reconstructed HDR image, by assuming exponentially distributed irradiance values. We define the weighting function as the probability that the observed pixel intensity is *approximately unbiased*. The weighting function can be directly computed based on the noise model parameters, which gives rise to different symmetric and asymmetric shapes when electronic noise or photon noise is dominant. We also explain how to deal with the case that some of the noise model parameters are *unknown* and explain how the camera response function can be estimated using the presented noise model. Finally, experimental results are provided to support our findings.

Introduction

The modeling of realistic camera noise is a subject that has not extensively been investigated, compared to the overwhelming amount of attention that the problem of stationary additive white Gaussian noise (AWGN) removal from images receives. A white stationary Gaussian noise assumption leads to simple and elegant denoising methods (such as wavelet shrinkage methods [1-5], total variation [6], anisotropic diffusion [7,8], NLMMeans [9-11]). However, when applied to realistic problems (e.g., the suppression of noise from CCD/CMOS measures taken with mobile phones or other consumer cameras), these techniques often yield poor results [11-13]. The main problem is the noise model mismatch, which causes the

techniques to either over- or underestimate the noise level. In realistic circumstances, noise is not an additive Gaussian process, not white nor stationary. For example, due to quantum-mechanical aspects inherent to the measurement of light, the sensor signals, in the form of analog voltages, are subject to statistical fluctuations with variances proportional to the “ideal”^a signal (i.e., the one we would like to measure). This phenomenon is known as *photon noise*. Subsequently, the measured analog voltages are amplified and converted to digital signals, leading to the introduction of *electronic noise* and *quantization noise*. The overall noise is a signal-dependent mix of noise from different sources, in combination with different linear and nonlinear pre/post-processing steps performed inside the camera. Consequently, the noise can not be well described using a stationary AWGN model.

*Correspondence: bart.goossens@telin.ugent.be
Department of Telecommunications and Information Processing
(TELIN/IP/iMinds), Ghent University, Sint-Pietersnieuwstraat 41, 9000 Gent,
Belgium

In literature, it is often advocated to use so-called *variance stabilizing transforms* [14,15], these are transforms that make the noise approximately “independent” and additive with a constant variance. The simple and elegant denoising approaches can then efficiently be applied to the variance stabilized signal and finally the inverse stabilization transform is applied. A first problem is the fact that a signal model that works fine in a linear domain may function poorly after variance stabilization [13]. A second problem is that *an accurate noise model* is needed in order to build such a variance stabilization transform.

Recently, a number of (simplified) CCD/CMOS noise models and estimation techniques have been presented in [12,16,17]. Next to this, many researchers address individual aspects of the noise in digital camera images (such as dealing with clipping [18], Poisson/multiplicative noise [19,20], signal-dependent/non-stationary noise [21-23]). Compared to these efforts, relatively little attention has gone to the *joint* modeling of the many involved factors that contribute to the noise characteristics in a digital camera. Some notable exceptions are [17,24,25]: Liu et al. [17] propose Bayesian estimation of the noise level function, directly from the camera reconstructed image itself. The set of possible noise level functions is thereby derived by considering the processing operations in the camera. In [24], a parametric noise model for RAW sensor data is proposed, which is then used to estimate the optimal exposure times for high dynamic range (HDR) image acquisition. A similar approach is presented in [25] for the photon-noise limited case.

In our opinion, image reconstruction techniques based on joint noise modeling techniques will (in the long term) significantly further improve the image quality, especially when physical limits of the image sensors are reached. For example, the minimal sensor size of a digital camera is limited by signal-to-noise considerations: if one would like to have a higher image resolution, one would also have to deal with a higher level of noise in the image, even up to such a degree that further increasing the image resolution does not bring any gain in image quality. Besides this, the dynamic range of a camera sensor is still many orders lower than the dynamic range that the human visual system can deal with, and improving the dynamic range will also bring an additional image quality increase.

Nowadays, to deal with these problems, camera manufacturers integrate extra post-processing operations, such as noise removal or HDR reconstruction into the camera. Their noise removal schemes are, mostly because of power consumption, hardware complexity and cost reasons, often based on the simple Gaussian noise models. Building more sophisticated and realistic noise models bridges the gap between the more and more miniaturized camera sensor designs and the increasing image quality expectations of camera end users.

Now that our goals are made clear, it is obvious that there are a lot of aspects (e.g. image noise, resolution, dynamic range, processing artifacts, ...) to cover in order to build such an accurate and realistic noise model. To limit the scope of this article, firstly we will assume a point-wise relationship between the input and the output, which can be described by a so-called camera (or intensity or brightness) response function (ignoring correlations both spatially and between different color channels). Secondly, we will limit ourselves to image artifacts that are caused by noise: we will not deal with other artifacts common in digital camera, such as chromatic aberrations, lens deformations, lens flares etc. Multivariate extensions of our theory and dealing with such artifacts may be a topic of future research. Also, while building sophisticated models, it is important to keep the models as simple as possible, so that techniques based on these models are still practical.

As an elegant illustration of our noise model, we will consider the problem of the reconstruction of a HDR image from a set of low dynamic range (LDR) images where the intensity range of each image covers a different part of the whole dynamic range (technically speaking, these images are captured by using different exposure times). This problem is illustrated in Figure 1. The application to HDR reconstruction is a beautiful example where many key image quality factors meet each other: first, we have the dynamic range that is limited by the camera sensor (and analog-to-digital converters), causing over- and underexposure in the images. Second, we are dealing with image noise, which is most apparent in images taken at low exposure times (or dark environments). Third, the “true” image resolution of the final HDR image is confined by the amount of image noise in the reconstructed HDR image. Fourth, color fidelity (i.e., whether the colors in the HDR image correspond well to the colors in the real scene) is an important image quality factor.

An important concept in HDR reconstruction is the *weighting function* (also called *certainty function*), which is introduced in [27] and subsequently used in [28-30]. The weighting function in [30] is defined as being proportional to the slope of the camera response function (CRF), indicating how quickly the output pixel intensity of a camera varies for a given input intensity. Correspondingly, certainty images are computed by applying certainty functions to digital images, revealing which parts of the image are the most “reliable”. It is found that this is the case for the midtones of the image. One of the issues however, is that several choices for weighting functions have been proposed in literature and that it is not clear which function to choose under which conditions.

In this article, we will formalize the concept of the “certainty function” in terms of a realistic camera model, and we will show that defining the certainty function as

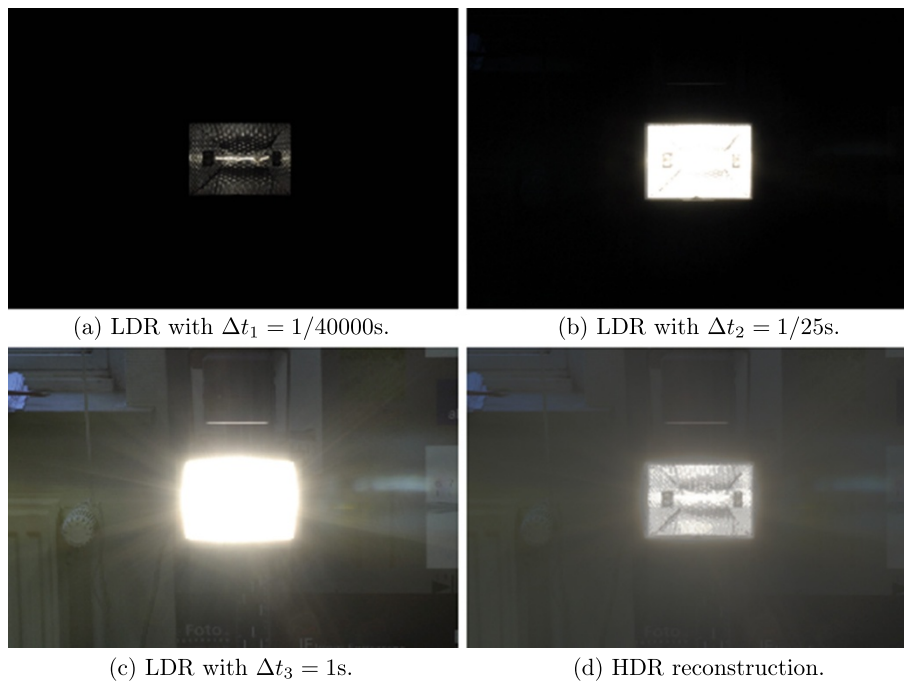


Figure 1 HDR Reconstruction Example. (a)–(c) Three out of nine LDR images of a halogen lamp taken with a Nikon D90 camera at different exposure times. (d) Synthesized HDR image (displayed with tone mapping from [26]), with details of the lamp and the background well reconstructed. Note that the LDR image in (c) contains contrast reducing lens flares and diffraction artifacts (which fall outside the scope of this article).

the probability that the output pixel intensity is *unbiased*, yields close to optimal HDR reconstruction results in the mean square error sense. The characteristics of our certainty function are similar to [30], additionally the Bayesian probabilistic formulation of our certainty function permits a more straightforward application to other problems, as we will demonstrate. To arrive at these novel certainty functions, we will develop a realistic camera noise model. Our noise model will be based on similar considerations as in [24], however the main difference is that we explicitly compute the bias functions based on an exponential prior distribution for the irradiance values, whereas in [24], simple indicator functions are used to predict whether the signal has been clipped.

The remainder of this article is organized as follows: in Section “Reconstruction of HDR images”, we introduce the synthesis problem of HDR images. Our realistic camera noise model is presented in Section “Camera noise modeling”. In Section “Probabilistic formulation of the weighting function”, we use the camera model to derive the weighting function. In Section “Estimation of the CRF”, we use the obtained weighting function in order to estimate the CRF for a set of LDR images. Results and a discussion are given in Section “Results and discussion”. Finally, Section “Conclusion” concludes this article.

Reconstruction of HDR images

Before going deep into the noise modeling for digital cameras, we first give an overview of the HDR reconstruction process and the different factors (such as SNR, the CRF and choice of the weighting function) that influence the reconstruction quality. First, we will introduce a number of general concepts that will be used throughout the article. Next, we will explain a HDR reconstruction technique that is based on a quite general noise model. Finally, we will investigate the “denoising” performance in the SNR sense of such a scheme, which will give some insight in the different factors that play a role in the quality of the final HDR image and will give an indication on conditions needed to obtain a certain minimal level of image quality, in terms of SNR.

Basic concepts

We have a set of $j = 1, \dots, P$ LDR digital photographs of a static scene at our disposal. The photographs are all taken from a fixed position, for example, using a tripod. We assume that lighting changes can be ignored, such that the incident scene radiance is constant during the exposure. For every photograph we use a different exposure time Δt_j , which will let us recover the dynamic range of the scene. Let z_{ij} denote the pixel intensity of photograph j at position i , where we use a one-dimensional index to

denote the spatial position (as in raster scanning). The goal of HDR reconstruction is to recover the irradiance map of the scene E_i , based on the pixel intensities z_{ij} of the digital photographs.

In a deterministic (noise-free) world, the image irradiance can be related to the pixel intensity of each photograph by the introduction of a CRF, as follows:

$$z_{ij} = \gamma \left(\Delta t_j \sqrt{\alpha_j} E_i \right) \quad (1)$$

where the point-wise function $\gamma(\cdot)$ is the CRF and with $\sqrt{\alpha_j}$ a gain factor (the squareroot is used here, for notational convenience later). The CRF models several digital post-processing operations performed in the camera after A/D conversion, such as: white balance, color corrections, demosaicing, gamma correction, edge/contrast enhancement, ... In Figure 2, an example of the CRF is given for a Nikon D90 camera. In this example, the CRF was estimated from the set of LDR images from Figure 1. This estimation technique will be discussed in more detail in Section "Estimation of the CRF". In practice, the CRF varies from camera to camera and even depends on (some of) the settings of a particular camera. The function is assumed to be point-wise (ignoring spatial information from the local neighborhood), mainly to keep the camera modeling and estimation simple. This has the drawback that some post-processing operations (such as demosaicing, sharpening, compression, color corrections, which can not be expressed in terms of a point-wise function), can not be included in the CRF. For the forthcoming analysis, we will first assume that the LDR images are saved in a RAW image format (which is possible for most high-end camera models). Thereby, the noise modeling becomes significantly easier, since post-processing operations, which severely affect the noise characteristics, do

not need to be taken into account. Later on, we will show the robustness of our approach, despite this assumption, on JPEG compressed LDR images where standard post-processing operations have been applied (see Section "HDR reconstruction from JPEG compressed LDR digital photographs"). Additionally, we assume that the CRF is *monotonically increasing* and hence invertible. This reflects the simple fact that relative positive changes in image irradiance at a certain position (e.g., due to illumination) should result in a positive change in the pixel intensity at the same position.

In real-life, there is not a one-to-one mapping of E_i to z_{ij} , because E_i is the *ideal* noise-free image irradiance, while z_{ij} is subject to noise. In this case, the CRF is the function that maps the *measured* image irradiance \hat{E}_i onto the pixel intensity z_{ij} , or mathematically speaking, $z_{ij} = \gamma \left(\Delta t_j \sqrt{\alpha_j} \hat{E}_i \right)$.

The concept of the CRF function is quite general: the techniques based on the CRF can also be used when all post-processing operations are switched off, for example for processing RAW sensor data. Because most camera manufacturers do not publicly share the internal processing steps of their cameras, together with their used parameter values, the CRF is usually estimated from the photographs themselves [28,30]. Two photographs with different exposure times are sufficient for this task.

According to (1), the *image irradiance* can be recovered by applying the inverse CRF to the pixel intensity and subsequently by dividing by the expose time. In the logarithmic domain, this can be expressed as (see [28], Eq. 2):

$$\log E_i = g(z_{ij}) - \log \Delta t_j - \frac{1}{2} \log \alpha_j, \quad (2)$$

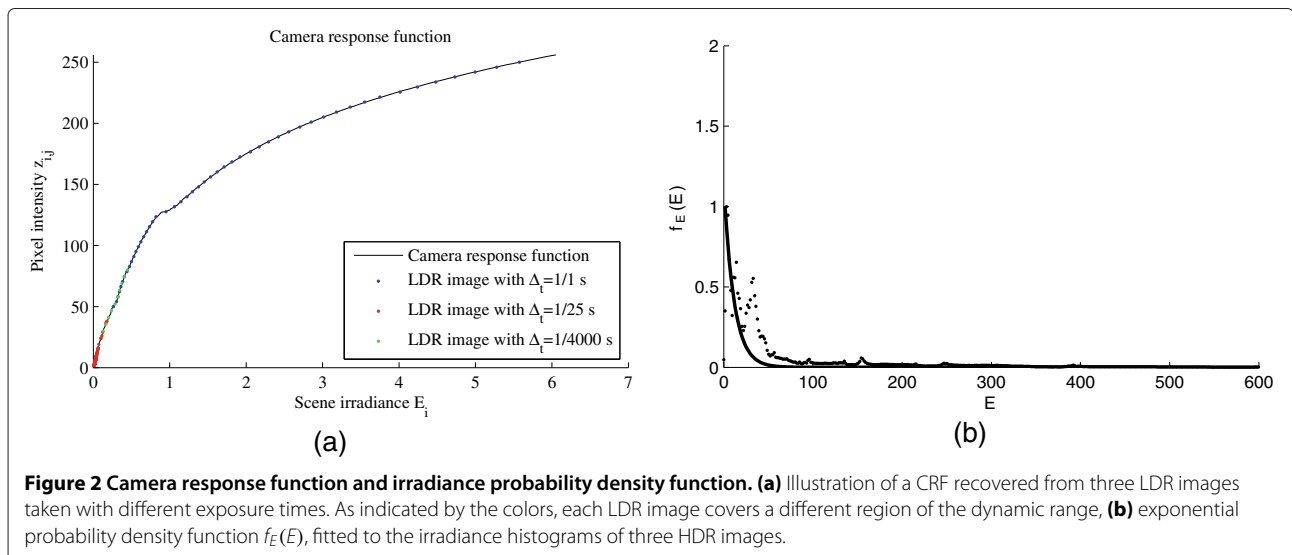


Figure 2 Camera response function and irradiance probability density function. (a) Illustration of a CRF recovered from three LDR images taken with different exposure times. As indicated by the colors, each LDR image covers a different region of the dynamic range, (b) exponential probability density function $f_E(E)$, fitted to the irradiance histograms of three HDR images.

with $g(z) = \log(\gamma^{-1}(z))$, the logarithm of the inverse CRF. Working in the logarithmic domain offers a number of practical advantages, such as the fact that estimates are linear in this domain. For example, writing (2) for several values of j gives a linear system of equations (in $\log E_i$ and $\log \Delta t_j$). To reconstruct a HDR image from a set of LDR images, the pixel intensities in two different LDR images can be related to each other as follows:

$$g(z_{ij}) - g(z_{ij'}) = \log \frac{\Delta t_j}{\Delta t_{j'}} + \frac{1}{2} \log \frac{\alpha_j}{\alpha_{j'}}, \quad (j \neq j'). \quad (3)$$

As an illustration, we depicted the intensities $g(z_{ij})$ as a function of the image irradiance E_i in Figure 3a, for three real LDR images taken with increasing exposure times Δt_j and equal gains α_j . Due to (2) and the different exposure times, there are three lines with different offsets. The difference in offsets can be computed using (3). Recovering an image with an extended dynamic range basically comes down to setting the offsets to zero (Figure 3b).

Noise model-based HDR reconstruction

In practice, the image irradiance measurements E_i are subject to statistical fluctuations called quantum or photon noise, caused by uncertainty associated by “counting” light energy quanta. Let $\hat{E}_i \Delta t_j$ denote an energy measurement after an exposure of Δt_j secs. The final pixel intensity of the LDR image j at position i is then given by $z_{ij} = \gamma(\hat{E}_i \sqrt{\alpha_j} \Delta t_j)$. Because of the statistical nature of \hat{E}_i , z_{ij} will be a random variable as well. The probability density function of z_{ij} is a complicated function in general, because:

- The measurement \hat{E}_i is affected by several (both additive and multiplicative) sources of noise, as we

will discuss in more detail in Section “Camera noise modeling”.

- A priori, little is known about the CRF $\gamma(\cdot)$ (except for the assumption of monotonicity).

Our solution to these problems is to firstly characterize the distribution of $\hat{E}_i \Delta t_j$ (or equivalently $g(z_{ij})$, if we exploit the fact that the CRF $\gamma(\cdot)$ is invertible, see Section “Basic concepts”). Secondly, we compute the distribution of z_{ij} through a change of variables. In Section “Camera noise modeling” we will explain that, under quite general circumstances, $g(z_{ij})$, conditioned on the log-irradiance, is well represented by the Gaussian distribution:

$$g(z_{ij}) | \log E_i \sim N(\log \sqrt{\alpha_j} \Delta t_j + \log E_i + v(E_i, \Delta t_j), \sigma^2(E_i, \Delta t_j)), \quad (4)$$

with $v(E_i, \Delta t_j)$ a bias term and with $\sigma^2(E_i, \Delta t_j)$ an irradiance and exposure time dependent variance (also see Figure 4). In Figure 5a, a probability density plot of $g(z_{ij}) | \log E_i$ is given. The point clouds correspond to histograms obtained after Monte-Carlo simulation, while the solid line is the approximated PDF in (4). Here, the asymmetry of the point cloud is mostly caused by taking the *logarithm* of the sum of a Poissonian and Gaussian distributed random variables. The asymmetry cannot be modeled well using the Gaussian distribution, however, in this work only the first two statistical moments (mean and variance) are of most importance.

The conditional Gaussianity of the PDF allows us to make a number of considerations, as we will show next. The HDR image can directly be reconstructed using the maximum likelihood method^b for estimating the mean

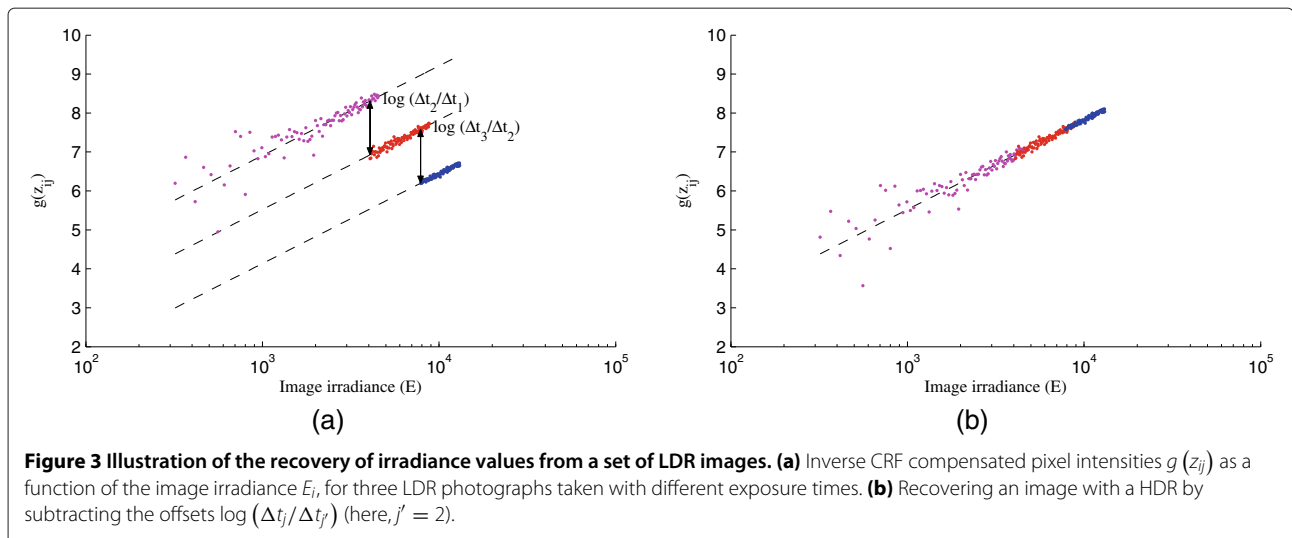
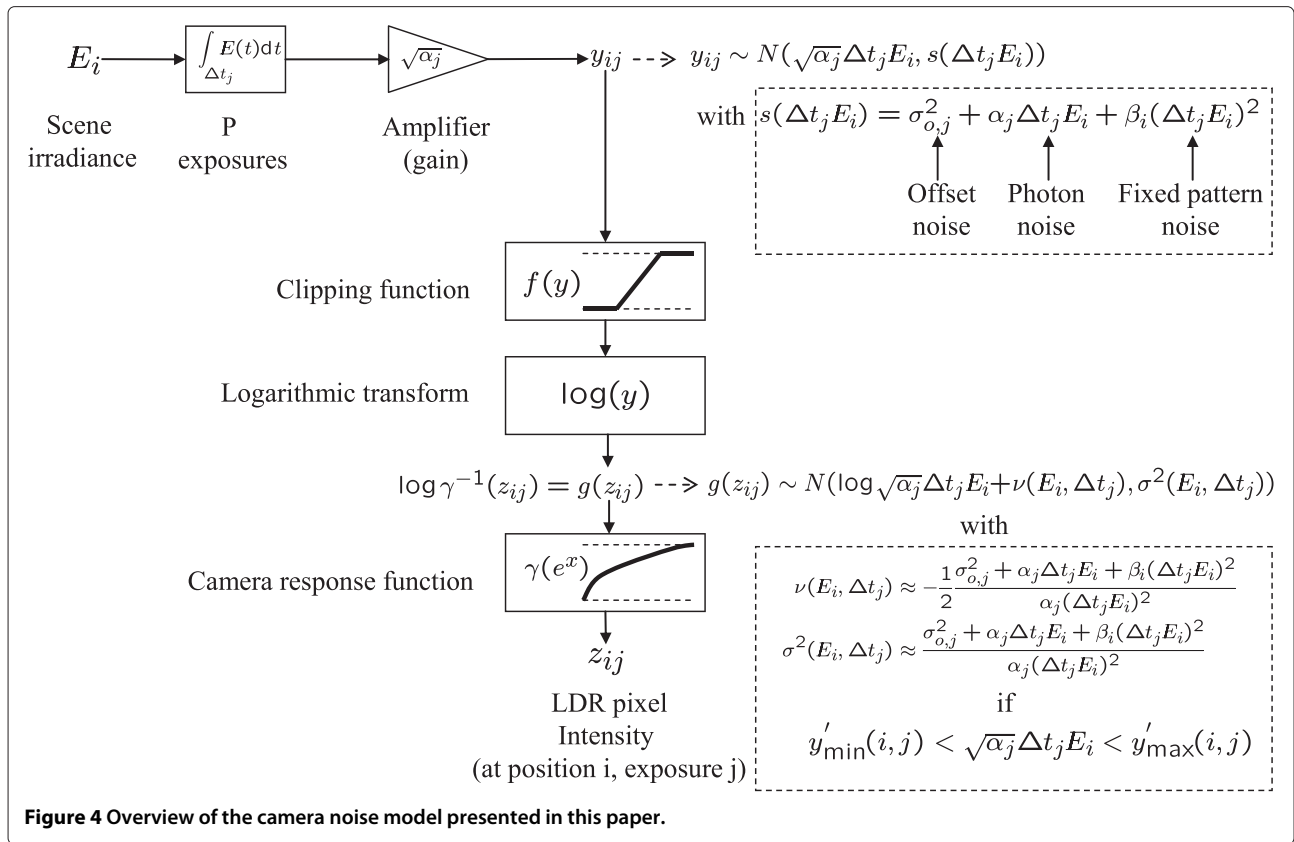


Figure 3 Illustration of the recovery of irradiance values from a set of LDR images. **(a)** Inverse CRF compensated pixel intensities $g(z_{ij})$ as a function of the image irradiance E_i , for three LDR photographs taken with different exposure times. **(b)** Recovering an image with a HDR by subtracting the offsets $\log(\Delta t_j / \Delta t_{j'})$ (here, $j' = 2$).

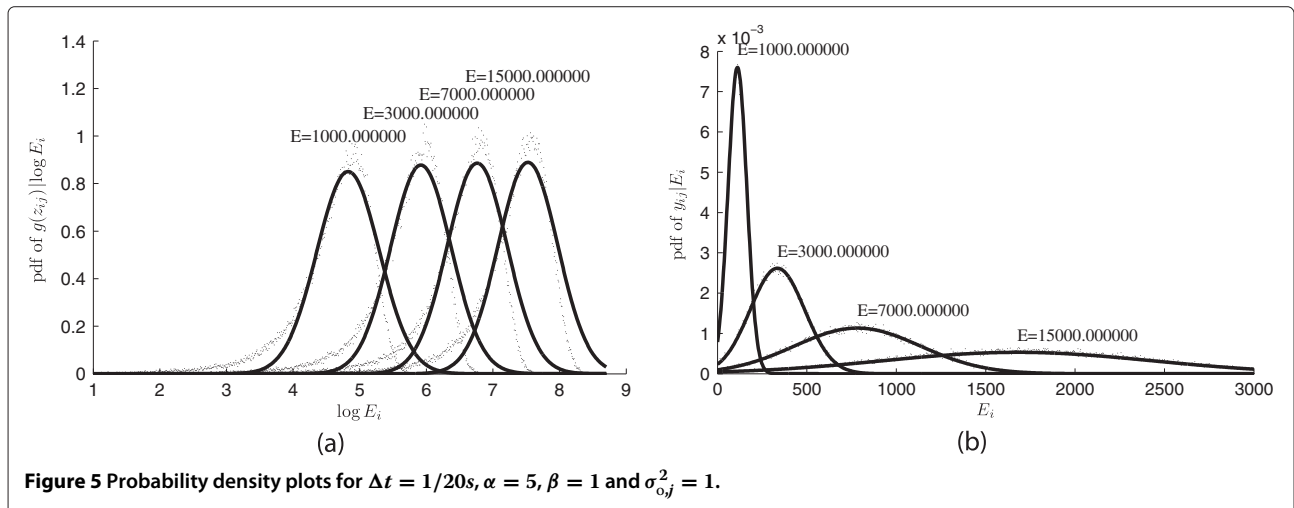


of a Gaussian distribution (i.e., by maximizing the log-likelihood function of the data):

$$\widehat{\log E_i} = \arg \max_{\log E_i} - \sum_{j=1}^P \frac{|g(z_{ij}) - \log \sqrt{\alpha_j} \Delta t_j - \log E_i - \nu(E_i, \Delta t_j)|^2}{\sigma^2(E_i, \Delta t_j)} \quad (5)$$

The factor $1/\sigma^2(E_i, \Delta t_j)$ takes into account that the noise variance (and correspondingly the SNR) is both exposure time and irradiance dependent, and consequently penalizes images with a high noise variance (low SNR) in the HDR reconstruction.

In practice, the “true” bias $\nu(E_i, \Delta t_j)$ is unknown and difficult to estimate from the LDR image data, due to the availability of a limited number of LDR images (typically, $P < 10$). We will explain later in Section “Camera noise



modeling” that, for z_{ij} in the middle of the camera’s exposure range (e.g., $z_{ij} = 128$ for 8-bit data), the bias is negligibly small compared to the image irradiance ($|\nu(E_i, \Delta t_j)| \ll \log \Delta t_j + \log E_i$) and hence can be ignored. A simple but different way to deal with the bias problem is to omit the term $\nu(E_i, \Delta t_j)$ and to introduce a weighting function $w_{ij} = w(z_{ij})$ as in [27-30]:

$$\widehat{\log E_i} = \arg \min_{\log E_i} \sum_{j=1}^P \frac{w_{ij}}{\sigma^2(E_i, \Delta t_j)} \quad (6)$$

$$|g(z_{ij}) - \log \sqrt{\alpha_j} \Delta t_j|^2.$$

A solution to this weighted least squares problem (6) is then given by:

$$\widehat{\log E_i} = \frac{\sum_{j=1}^P w_{ij} \sigma^{-2}(E_i, \Delta t_j) (g(z_{ij}) - \log \sqrt{\alpha_j} \Delta t_j)}{\sum_{j=1}^P w_{ij} \sigma^{-2}(E_i, \Delta t_j)}, \quad (7)$$

which we will call the main HDR reconstruction formula. The estimated log-irradiance is a weighted average of $g(z_{ij}) - \log \Delta t_j$, where each term can be interpreted as an estimate of the log-irradiance for each individual LDR image. A schematic overview of this reconstruction method is given in Figure 6.

Let us now turn to the choice of the weighting function, which is crucial for the correct operation of the HDR algorithm, since the weighting function determines the

trade-off between dynamic range and SNR. In literature, several functions have been proposed:

- Debevec and Malik [28] propose a triangular function to emphasize the middle of the exposure range, mainly for its simplicity:

$$w_{ij}(z) = \begin{cases} z - z_{\min} & z \leq z_{\text{mid}} \\ z_{\max} - z & z > z_{\text{mid}} \end{cases} \quad (8)$$

with $z_{\text{mid}} = \frac{1}{2} (z_{\min} + z_{\max})$. We will see in Section “Dealing with unknown camera model parameters” that this weighting function is a good choice when little information is known about the camera noise characteristics.

- Mann and Picard [27,30] select a weight related to the slope of the CRF γ , which indicates how quickly the pixel intensity z_{ij} varies with the input, in order to assign lower weights to coarser quantized pixel intensities. It is argued that for noisy input images, influences of quantization noise are minimal in the middle of the camera’s exposure range. Using our notations, the weighting function is defined as follows ([30], Eq. 11):

$$w_{ij}(z) = \frac{dz}{d \log \hat{q}_j(z)} = \hat{q}_j(z) \frac{d\gamma}{dz} (\hat{q}_j(z)) \quad (9)$$

where $\hat{q}_j(z) = \gamma^{-1}(z) / (\Delta t_j \sqrt{\alpha_j})$ is an estimate for the image irradiance, based on exposure j . For an identity CRF, the weights are approximately proportional to the image irradiance.

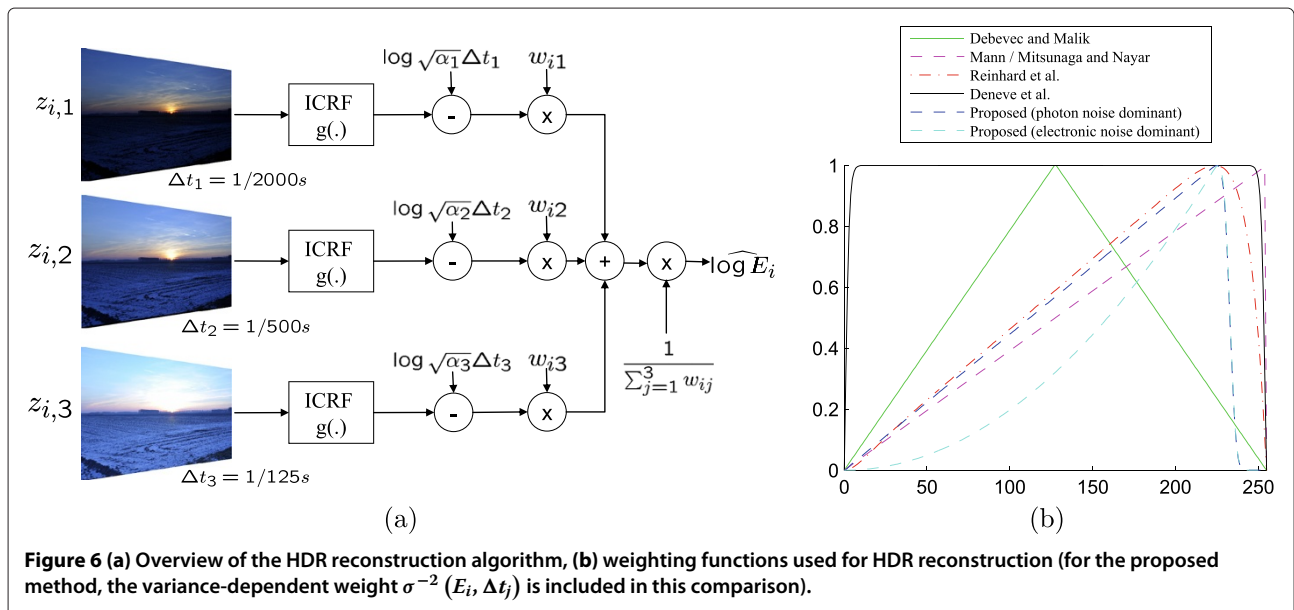


Figure 6 (a) Overview of the HDR reconstruction algorithm, (b) weighting functions used for HDR reconstruction (for the proposed method, the variance-dependent weight $\sigma^{-2}(E_i, \Delta t_j)$ is included in this comparison).

- Mitsunaga and Nayar [29] relate the weights to the SNR of the LDR image (which is assumed to be linear to the image irradiance). Thereby, the authors assume that the measurement noise is both stationary and independent of the underlying signal. This results in the following equation:

$$w_{ij}(z) = \hat{q}_j(z) / \frac{d\hat{q}_j(z)}{dz} (z) = \hat{q}_j(z) \frac{d\gamma}{dz} (\hat{q}_j(z)), \quad (10)$$

where the rule for the derivative of the inverse was applied. The weighting function proposed by Mitsunaga and Nayar is the same as the weighting function of Mann.

- In previous work [31], we proposed an exponential power function as a trade-off between noise suppression and clipping:

$$w_{ij}(z) = \exp\left(-\frac{(z - z_{\text{mid}})}{b(z_{\text{mid}})^a}\right), \quad (11)$$

where a and b are constants which were determined experimentally. This function also emphasizes the middle of the exposure range, but in such a way that the low and high intensities are still assigned a large weight, close to 1.

- Other weighting functions are proposed by Reinhard et al. [32], Tsini et al. [33], Kirk and Andersen [34].

We remark that in many of the above works, the authors assume AWGN that is independent of the image irradiance and image index j , which leads to a special case of (6) in which $\sigma^2(E_i, \Delta t_j)$ is constant (consequently this factor can be dropped in (7)). In general, because of the dependency of the weighting factor $\sigma^{-2}(E_i, \Delta t_j)$ on the *unknown* image irradiance E_i , $\widehat{\log E_i}$ needs to be estimated as in an iteratively re-weighted least squares approach. During this iterative process, in the right handed side of (7), the image irradiance estimate E_i of the previous iteration is used. The image irradiance is then estimated using the currently best available estimate of the weight $w_{ij}\sigma^{-2}(E_i, \Delta t_j)$. Once a better estimate E_i becomes available, the combined weights $w_{ij}\sigma^{-2}(E_i, \Delta t_j)$ are updated.

Such an iterative updating scheme can be avoided (which is advantageous from a computational point of view), even *without* assuming AWGN: if $\sigma^{-2}(E_i, \Delta t_j)$ is proportional to $(E_i)^q$, the factors $(E_i)^q$, which are independent of the summing variable j , in the numerator and denominator of (7) cancel each other. In the next section, we will perform a thorough modeling of the camera noise characteristics, to obtain explicit formulas for $\sigma^{-2}(E_i, \Delta t_j)$. Surprisingly, our result in Section “Camera noise modeling” indicates that an approximation such as $\sigma^{-2}(E_i, \Delta t_j) \propto (\Delta t_j E_i)^q$ with $q = 1$ or 2 are quite adequate for a digital camera noise model.

Denoising performance of the HDR reconstruction formula

As we explained before, the final reconstructed HDR image is a weighted average of the different LDR images, with a correction term to compensate the different exposure times. Because of the averaging (which increases the SNR), denoising is intrinsically part of HDR reconstruction. To gain some insight in the different factors that determine the quality of the reconstructed HDR image, it is useful to investigate the “denoising performance” of the main HDR reconstruction formula (7). To start, we will perform our following analysis for a constant weighting function $w_{ij} = 1, j = 1 \dots, P$. Based on (7), we can write the variance of $\widehat{\log E_i}$ as:

$$\begin{aligned} \text{Var} \left[\widehat{\log E_i} \right] &= \frac{\sum_{j=1}^P \sigma^{-4}(E_i, \Delta t_j) \text{Var} [g(z_{ij})]}{\left[\sum_{j=1}^P \sigma^{-2}(E_i, \Delta t_j) \right]^2} \\ &= \left[\sum_{j=1}^P \sigma^{-2}(E_i, \Delta t_j) \right]^{-1}. \end{aligned} \quad (12)$$

Exploiting the fact that the SNR is defined in terms of the ratio of the signal energy ($(\log E_i)^2$) and the noise energy ($\text{Var} [\widehat{\log E_i}]$), the signal-to-noise ratio of the reconstructed HDR image can be expressed in an elegant formula:

$$\begin{aligned} \text{SNR}_{\text{HDR}} &= 10 \log_{10} \left((\log E_i)^2 / \text{Var} [\widehat{\log E_i}] \right) \\ &= 20 \log_{10} \log E_i \\ &\quad + 10 \log_{10} \sum_{j=1}^P \sigma^{-2}(E_i, \Delta t_j) \text{ [dB]}. \end{aligned} \quad (13)$$

We will later see, that, under some conditions (more specifically, when photon noise is dominant), $\sigma^{-2}(E_i, \Delta t_j) \approx E_i \Delta t_j$, which allows us to approximate the SNR as:

$$\text{SNR}_{\text{HDR}} \approx 20 \log_{10} (E_i \log E_i) + 10 \log_{10} \sum_{j=1}^P \Delta t_j \text{ [dB]}. \quad (14)$$

such that each extra LDR image increases the SNR by maximally^c $10 \log_{10} \Delta t_j$ dB (see Figure 7). For this model, three factors can increase the SNR: (1) increasing the light incidence of the scene (for example by example by turning on an additional light source), (2) increasing the exposure Δt_j and (3) increasing the exposures P . Obviously, some of these measures come at a cost: for example, when increasing the exposure or the number of exposures, the scene stationary assumption may not hold any longer due to object motion. However, the formula $\text{SNR}_{\text{HDR}} > 0$ dB expresses necessary (although not

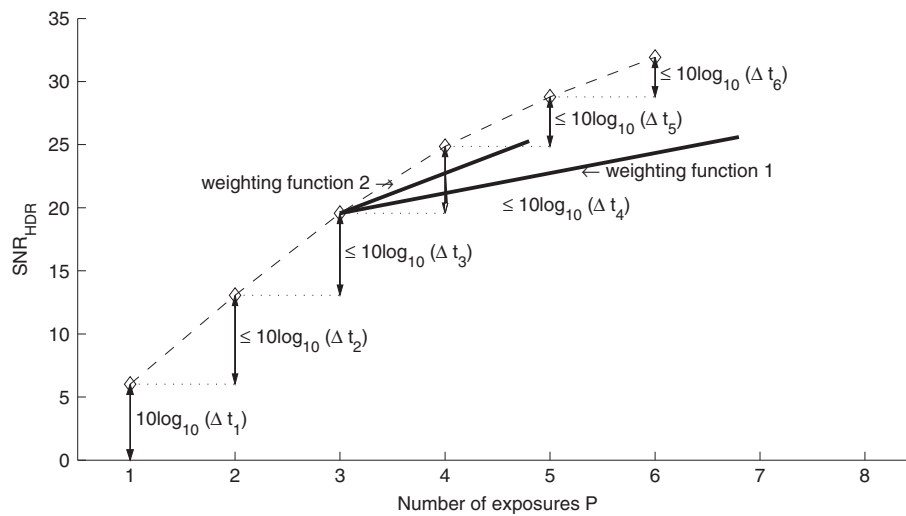


Figure 7 Maximal theoretical contribution of each LDR image in the overall SNR of the reconstructed HDR image and the influence of the weighting function.

sufficient) conditions required for practical HDR reconstruction. HDR reconstruction is not appropriate if the SNR reaches zero ($SNR_{HDR} < 0$ dB, or equivalently $\sum_{j=1}^P \Delta t_j < (E_i \log E_i)^2$). In this case, the so-called Wyck-off signal-noise criterion [30] is violated: the exposure is not high enough to overcome the sensor noise.

In case the weighting function is non-constant, the variance becomes:

$$\begin{aligned} \text{Var} [\widehat{\log E_i}] &= \frac{\sum_{j=1}^P \sigma^{-2} (E_i, \Delta t_j) w_{ij}^2}{\left[\sum_{j=1}^P \sigma^{-2} (E_i, \Delta t_j) w_{ij} \right]^2} \\ &\geq \left[\sum_{j=1}^P \sigma^{-2} (E_i, \Delta t_j) \right]^{-1}. \end{aligned} \quad (15)$$

In other words, non-constant weighting function, such as the weighting functions proposed in literature, lead to a reduced SNR_{HDR} . Figure 7 also illustrates how the SNR of the reconstructed HDR image is affected by using different

weighting functions. We will go deeper into this topic in Section “Probabilistic formulation of the weighting function”.

Camera noise modeling

In Section “Noise model-based HDR reconstruction” we have put forward a Gaussian model for the inversely compensated pixel intensity $g(z_{ij})$. This Gaussian model comprises next to the mean $\log(\Delta t_j E_i)$, a bias function $v(E_i, \Delta t_j)$ and a noise variance function $\sigma^2(E_i, \Delta t_j)$. For accurate HDR reconstruction, expressions for these functions are indispensable. These expressions through camera noise modeling. Our noise modeling and the way in which we deal with clipping is similar to the Poissonian–Gaussian modeling of [18,35], with the main difference that our analysis covers the more general case in which the CRF is nonlinear (see Appendix 1). But before building a camera noise model, we first need a better understanding of the processing inside a digital camera.

In Figure 8, a simplified block scheme is given of the processing pipeline of a digital camera. First, the incident light of the scene enters the camera sensor through

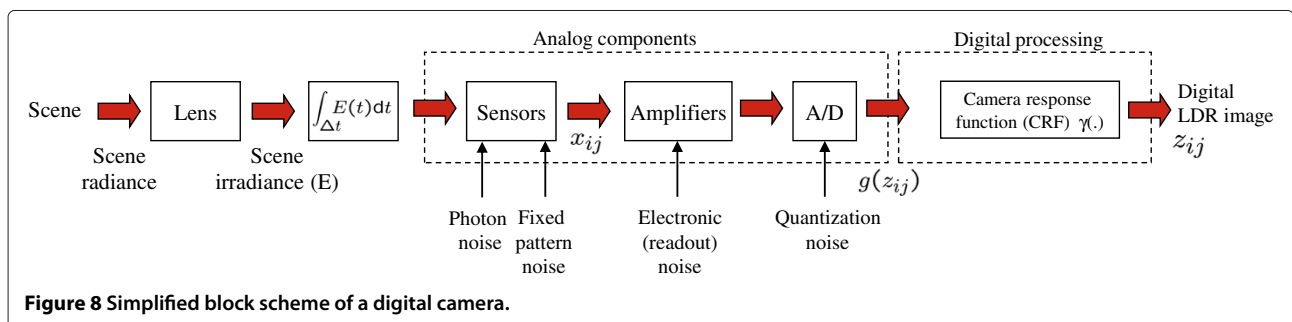


Figure 8 Simplified block scheme of a digital camera.

the lens. Because the image irradiance E_i is constant during the expose of length Δt_j , the energy registered by the sensor will then roughly be the product of the image irradiance and the exposure time. As already mentioned, the measured energy is subject to photon noise. The photon noise is mostly noticeable in the image when taking photographs in dark environments or at low exposure times. A commonly used probability density model for the sensor noise is the Poisson distribution:

$$x_{ij} \sim \mathcal{P}(\Delta t_j E_i) \quad (16)$$

where x_{ij} is the measured intensity by the photo-detector at position i . We will model the photon noise using a Gaussian distribution with the same first and second order moments as the Poisson distribution, i.e. $x_{ij} \sim \mathcal{N}(\Delta t_j E_i, \Delta t_j E_i)$, which is a good approximation if $\Delta t_j E_i \gg 0$. This is the case for consumer digital cameras [13] and also considerably simplifies the modeling task.^d In the next step, x_{ij} is amplified and quantized by the analog-to-digital converter (A/D) to B bits. More precisely, the analog voltages x_{ij} measured by the photo-detectors are subjected to the following operations:

1. *Amplification*, which introduces electronic noise. This electronic noise is well characterized by a Gaussian distribution $\mathcal{N}(0, \sigma_\epsilon^2)$. We will denote by $\sqrt{\alpha_j}$ the amplification gain for image j , resulting in a signal with expected value $\sqrt{\alpha_j} \Delta t_j E_i$. The amplification gain is usually determined by the ISO sensitivity of the camera.
2. The measured signal x_{ij} is affected by non-uniform heating of the sensor, and becomes non-uniform, even when the scene radiance is constant. The resulting fluctuations are called *fixed-pattern noise* (FPN), also known as dark current non-uniformity. FPN has a variance that is quadratic in the image irradiance (mainly caused by variations by variations in the gain factor α_j , see [36]), but also contributes to the overall noise that is independent of the image irradiance (as an offset term). We will call these two noise components respectively gain FPN and offset FPN.
3. *A/D conversion*, which involves clipping of the dynamic range to the range $0-2^B - 1$ and quantization. The camera sensor can only map irradiance values that are within certain minimum and maximum bounds (these bounds determine the *tonal range* of the sensor) and values that are outside this range are *clipped*. Because the tonal range of the sensor often cannot cover the whole dynamic range of the scene, some regions in the image will be over- or underexposed. In the

following $f(x) = \max(0, \min(2^B - 1, x))$ will denote this clipping function.

4. For ideal A/D converters, the *quantization noise* is uniformly distributed in the range $[-1/2, 1/2]$. In this work, we will treat quantization noise similarly to electronic noise and offset FPN. This is possible since quantization is performed directly after amplification. Therefore, the offset noise variance $\sigma_{o_j}^2$ signifies the summed contributions of the variances of the electronic noise, offset FPN and quantization noise.
5. Finally, the quantized signal is subjected to several *post-processing steps* in the digital domain. Examples are: gamma correction, brightness and contrast adjustment, color corrections, white balance, compression/expansion etc. As we mentioned before, these operations are modeled using the CRF.

The different processing steps outlined above significantly alter the statistical properties of the measured intensities in Equation (16). In this work we separate *linear* effects from *nonlinear* effects and model them separately. For example, while amplification (linearly) magnifies the intensity by a factor $\sqrt{\alpha_j}$, the variance of the intensity is multiplied by a factor α_j . Similarly, adding offset noise (a linear additive process), increases the variance by $\sigma_{o_j}^2$. The clipping function and the CRF are both nonlinear functions. Because of the order of the steps 1–5, the nonlinear functions are positioned at the back of the chain, which simplifies the modeling task. y_{ij} the signal after applying steps 1–4, then we have the statistical dependency $x_{ij} \rightarrow y_{ij} \rightarrow z_{ij}$. Using the Gaussian approximation of x_{ij} , the PDF of y_{ij} can readily be found (see Figure 4):

$$y_{ij} \sim \mathcal{N}\left(\sqrt{\alpha_j} \Delta t_j E_i, \sigma_{o_j}^2 + \alpha_j \Delta t_j E_i + \beta_{ij} (\Delta t_j E_i)^2\right), \quad (17)$$

where β_{ij} is a gain of FPN whose variance increases as a quadratic function of E_i (see step 2). To check the normality of y_{ij} , a normal probability plot is provided in Figure 5b. It can be noted that the Gaussian approximation is quite accurate, even for small irradiance values (e.g., $\Delta t_j E_i < 100$). Next, the measured pixel intensities z_{ij} are related to y_{ij} by the nonlinear relationship:

$$z_{ij} = \gamma(f(y_{ij})). \quad (18)$$

Now, the noise variance $\sigma^2(E_i, \Delta t_j)$ and bias $v(E_i, \Delta t_j)$ can be computed from the statistical moments of $\log f(y_{ij})$ using the definitions of mathematical expectation and variance, as shown in (19).

$$\begin{aligned}
 \nu(E_i, \Delta t_j) &= E[\log f(y_{ij}) | E_i] - \log f(\sqrt{\alpha_j} \Delta t_j E_i) \\
 &= C \int_{-\infty}^{+\infty} \log \frac{f(u) du}{f(\Delta t_j E_i)} \exp \\
 &\quad \times \left(-\frac{1}{2} \frac{(\sqrt{\alpha_j} \Delta t_j E_i - u)^2}{s(\Delta t_j E_i)} \right) \\
 \sigma^2(E_i, \Delta t_j) &= \text{Var}[\log f(y_{ij}) | E_i] \\
 &= C \int_{-\infty}^{+\infty} (\log f(u))^2 du \exp \\
 &\quad \times \left(-\frac{1}{2} \frac{(\sqrt{\alpha_j} \Delta t_j E_i - u)^2}{s(\Delta t_j E_i)} \right) \\
 &\quad - C \left[\int_{-\infty}^{+\infty} \log f(u) du \exp \right. \\
 &\quad \left. \times \left(-\frac{1}{2} \frac{(\sqrt{\alpha_j} \Delta t_j E_i - u)^2}{s(\Delta t_j E_i)} \right) \right]^2
 \end{aligned} \tag{19}$$

where $s(x) = \sigma_{o,j}^2 + \alpha_j x + \beta_{ij} x^2$ and $C = (2\pi s(\Delta t_j E_i))^{-1/2}$.

Although the functions in (19) can be calculated numerically, it is not directly clear how $\sigma^2(E_i, \Delta t_j)$ is related to the parameters $\sigma_{o,j}^2$, α_j , β_{ij} and Δt_j . Thereby it becomes very difficult to, e.g., devise techniques to estimate these parameters from the data z_{ij} .

In Appendix 1, we describe a technique that allows us to approximate the statistical moments of a general nonlinear function of y_{ij} , resulting in far simpler expressions. The premise is that the technique cannot be used when the nonlinear function is not differentiable. This is the case for the clipping function near its clipping points. If we disregard the clipping function (we will explain in a few moments how to deal with clipping), using Appendix 1, the statistical moments are well approximated by:

$$\nu(E_i, \Delta t_j) \approx -\frac{1}{2} \frac{\sigma_{o,j}^2 + \alpha_j \Delta t_j E_i}{\alpha_j \Delta t_j^2 E_i^2} - \frac{\beta_{ij}}{2\alpha_j} \quad \text{and} \tag{20}$$

$$\sigma^2(E_i, \Delta t_j) \approx \frac{\sigma_{o,j}^2 + \alpha_j \Delta t_j E_i}{\alpha_j \Delta t_j E_i} + \frac{\beta_{ij}}{\alpha_j}. \tag{21}$$

Interestingly, the gain FPN appears as an extra bias term in (20). Since FPN does not change from one image to another, a FPN template can be constructed for each ISO setting. Gain FPN removal then simply consists of subtracting the template $-\beta_{ij}/2\alpha_j$ from $\log f(y_{ij})$, the logarithm of the clipped digitized signal. This bias subtraction approach is quite common in digital cameras.

Let us now analyze the noise variance (21). In case the photon noise is dominant, i.e. $\alpha_j \Delta t_j E_i \gg \sigma_{o,j}^2$ and $\alpha_j \gg \Delta t_j E_i \beta_{ij}$, inversely proportional to $\Delta t_j E_i$. On the other hand, if the gain FPN is dominant, the noise variance is inversely proportional to $\Delta t_j^2 E_i^2$:

$$\sigma^2(E_i, \Delta t_j) \approx \begin{cases} \frac{1}{\Delta t_j E_i} & \text{(photon noise limited region)} \\ \frac{\sigma_{o,j}^2}{\alpha_j \Delta t_j^2 E_i^2} & \text{(offset noise limited region)} \\ \frac{\beta_{ij}}{\alpha_j} & \text{(gain FPN limited region)} \end{cases} \tag{22}$$

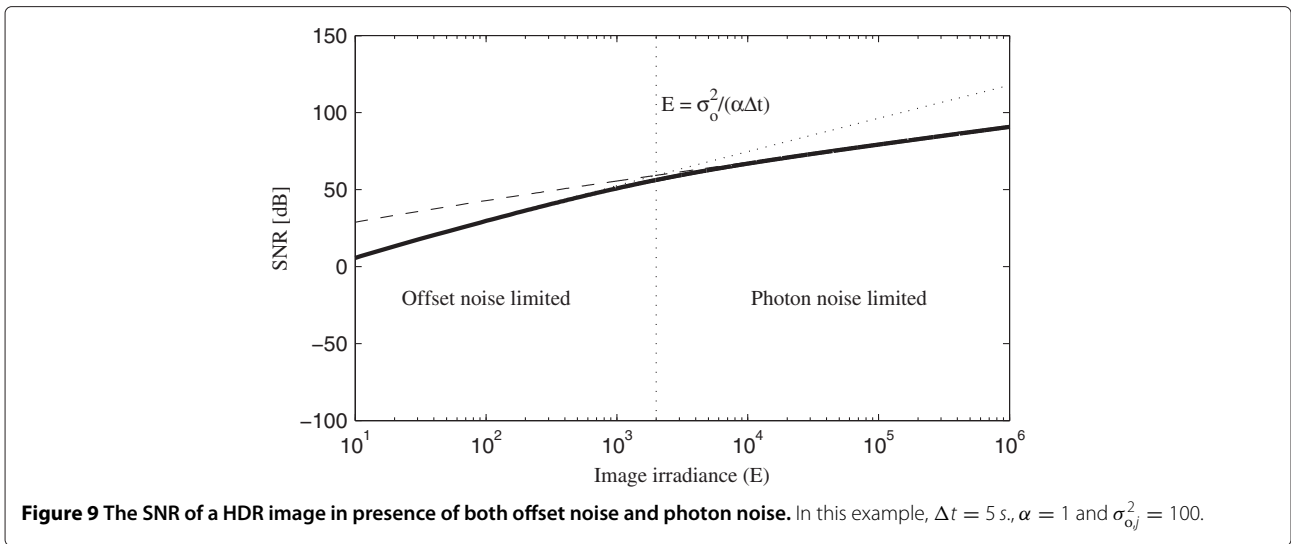
such that $\sigma^{-2}(E_i, \Delta t_j)$ can be approximated by a monomial in E_i . From Equation (22), the SNR of each image can easily be computed as $\text{SNR} = 20 \log_{10} \log(\sqrt{\alpha_j} \Delta t_j E_i) / \sigma(E_i, \Delta t_j)$. The SNR for the dominant photon and offset noise cases are equal for $E_i = \sigma_{o,j}^2 / (\alpha_j \Delta t_j)$, the SNR breaking point. In Figure 9, a plot of the SNR is given as a function of the image irradiance. While the “exact” SNR is a complicated nonlinear function in general (see (19)), the function can be well approximated by a piecewise linear function (in a logarithmic scale on both the horizontal and vertical axes). As explained at the end of Section “Noise model-based HDR reconstruction”, this finding allows us to develop efficient and practical HDR reconstruction algorithms based on a real camera noise model.

One of our main issues at this stage is that the approximations (20) and (21) only hold for differentiable CRFs, hence clipping of the dynamic range is not incorporated in our model. To resolve this issue, in Appendix 2, we derive bounds for $\Delta t_j E_i \sqrt{\alpha_j}$ for which the approximations (20) and (21) are accurate (up to an arbitrary precision). If we denote the minimal and maximal intensity values (after applying the clipping function) as, respectively y_{\min} and y_{\max} , we have shown that the results hold (up to an arbitrary precision) in a “clipping-free” region $y'_{\min} \leq \Delta t_j E_i \sqrt{\alpha_j} \leq y'_{\max}$, where typically $y'_{\min} < (y_{\min} + y_{\max}) / 2 < y'_{\max}$ (with y'_{\min} and y'_{\max} as in Appendix 2). An illustration is given in Figure 10, for $\Delta t = 1$ s, $\alpha = 1$ and $\sigma_{o,j}^2 = 0.01$.

Even though complicated processing by several nonlinear functions is involved, we can well describe the first and second statistical moments of the variables $g(z_{ij})$ using a convenient formula that holds with great accuracy within certain bounds. This finding will turn out to be very useful for applications that depend on this noise model. In the next sections, we will take these bounds into account in the HDR reconstruction, in particular for designing an appropriate weighting function.

Probabilistic formulation of the weighting function

The camera noise model from previous section can relatively easily be applied in many practical applications. The



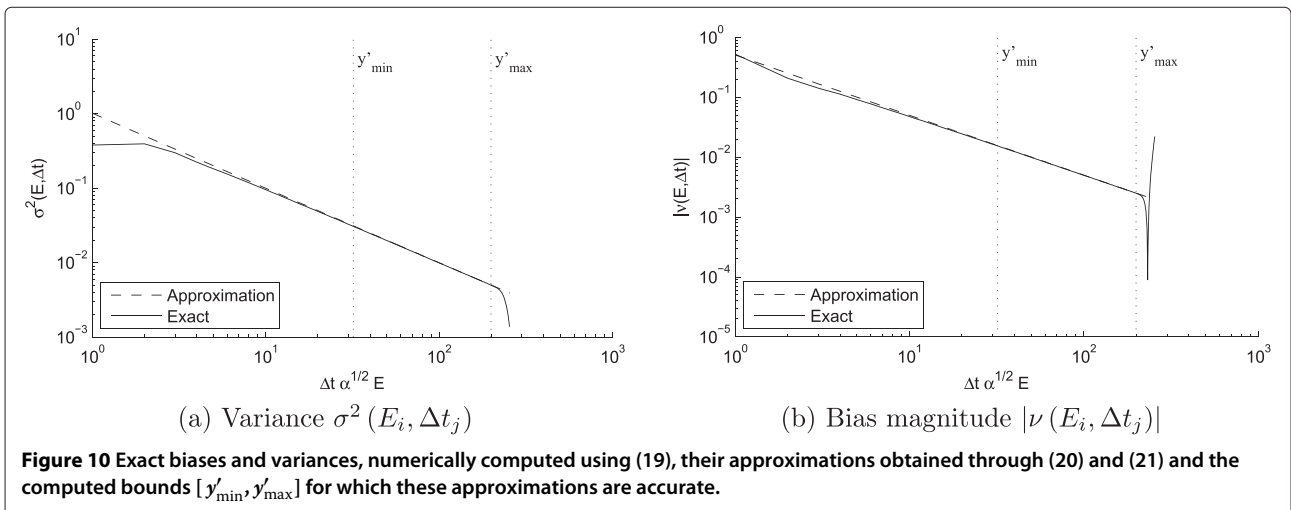
main ingredients of this model are the approximations of the noise variance function and the bias function, together with the ranges $[y'_{\min}, y'_{\max}]$ on which these approximations are accurate.

From the explanation in Section “Denoising performance of the HDR reconstruction formula”, the reader may expect that an optimal weighting function is one that is constant everywhere in the clipping-free regions ($w_{ij} = 1 - I(z_{ij} \leq 0) - I(z_{ij} \geq 2^B - 1)$, with $I(\cdot)$ the indicator function). However, such weighting function does *not* give maximal SNR because of bias effects: for example, an initially very large intensity y_{ij} may become less than $2^B - 1$ with a certain probability, due to addition of offset noise or FPN, but without being affected by the clipping operation. On average, several of these initially large intensities cause a *biased* HDR image estimate, when the weighting function is not properly chosen. In this section, we will

show how the camera noise model can be used to compute a near optimal weighting (or certainty) function in terms of SNR to be used in combination with the reconstruction formula (7).

MMSE-based estimation of the weighting function

Before deriving the weighting function, we make two observations on the requirements for the weighting function. As already mentioned, we would like to have an HDR reconstructed image that is *unbiased*: $E[\widehat{\log E_i}] = \log E_i$. Practically speaking, this condition means that the reconstructed log-irradiance is a true estimate of the “ideal” image log-irradiance. Because the brightness perception of the human eye is approximately *logarithmic* in the irradiance values (this is known as the Weber–Fechner law), we use the logarithm of the image irradiance. Using (7)



and (4), the condition of unbiased estimates is equivalent to:

$$E \left[\widehat{\log E_i} \right] = \log E_i + \frac{\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j) v (E_i, \Delta t_j)}{\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j)}. \quad (23)$$

Here we simply computed the mathematical expectation of both sides of (7). Note that a *sufficient* condition for an unbiased estimate is given by $v (E_i, \Delta t_j) \neq 0 \Rightarrow w_{ij} = 0$. Our approximation theory (Appendix 2) then states that the latter condition is the case if $y'_{\min} \leq \Delta t_j E_i \sqrt{\alpha_j} \leq y'_{\max}$.

Next to having unbiased estimates, the SNR of $\widehat{\log E_i}$ is an important factor, indicating the quality of the reconstructed HDR image. To take both effects into account, we will derive the Bayesian minimum mean square error (MMSE) estimator for our problem. We will minimize the MSE between the estimated log-irradiance and its true value, which is defined by:

$$\begin{aligned} \text{MSE} &= E \left[\left(\widehat{\log E_i} - \log E_i \right)^2 \right] \\ &= \underbrace{\text{Var} \left[\widehat{\log E_i} \right]}_{\text{variance term}} + \underbrace{\left(E \left[\widehat{\log E_i} - \log E_i \right] \right)^2}_{\text{squared bias term}} \end{aligned} \quad (24)$$

where we split up the MSE into two terms: the first term is closely related to the SNR of the reconstructed HDR image (see Equation (13)), the second term is the expected *squared bias error*. Next, we want to minimize the MSE with respect to the values of the weighting function: $\min_{w_{ij}} \text{MSE}$. The weighting function w_{ij} is positive $w_{ij} \geq 0$, and additionally we add the constraint $\max_i w_{ij} = 1$ to get rid of the scaling ambiguity.⁶ For the HDR reconstruction formula from Equation (7), the different terms of the MSE can be written as follows:

$$\text{Var} \left[\widehat{\log E_i} \right] = \frac{\sum_{j=1}^P \sigma^{-2} (E_i, \Delta t_j) w_{ij}^2}{\left[\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j) \right]^2}, \quad (25)$$

$$E \left[\widehat{\log E_i} - \log E_i \right] = \frac{\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j) v (E_i, \Delta t_j)}{\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j)}. \quad (26)$$

Note that both the variance and squared bias term have the same denominator, which will allow us to find an explicit linear solution to this minimization problem, as we will explain next. However, because taking the square of the bias term (26) results in several cross-terms $w_{ij} w_{ij'}$, which eventually leads to a linear system of equations with a non-diagonal system matrix, we will look for a solution in which these cross terms do not appear. This will have as a practical benefit that the weighting function for each LDR image can be calculated *solely* based on the model parameters of the particular image, which permits offline

computation of the weighting function. To achieve this mathematically, we will minimize an *upper bound* for the MSE. This upper bound should be selected carefully, close to the true MSE, such that the corresponding solution can be considered to be a good approximation to the solution of the original problem (Equation (24)). Using the inequality of Cauchy-Schwartz,⁷ such an upper bound can easily be found:

$$\text{MSE} \leq \frac{\sum_{j=1}^P [w_{ij} \sigma^{-2} (E_i, \Delta t_j)]^2 (\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j))}{\left(\sum_{j=1}^P w_{ij} \sigma^{-2} (E_i, \Delta t_j) \right)^2}. \quad (27)$$

In Appendix 3 the general solution for the minimization problem of (27) is derived. Using this result, we can write the optimal weighting function (minimizing (27) and relying on $\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j) > 0$)⁸ as:

$$w_{ij} = A \frac{\sigma^2 (E_i, \Delta t_j)}{\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j)} \quad (28)$$

with $A = [\max_j \sigma^2 (E_i, \Delta t_j) / (\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j))]^{-1}$ a normalization factor. According to (28), the optimal weights minimizing (27), can be found by computing the proportion of the noise variance $\sigma^2 (E_i, \Delta t_j)$ compared to the whole $\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j)$. We remark that when the bias $v (E_i, \Delta t_j) = 0$, the corresponding weight is equal to one. When the bias increases, the weight becomes smaller. Readers familiar with Wiener filters will note that (28) resembles the classical scalar Wiener filter weight formula, however (28) is not related to the Wiener filter because $P v^2 (E_i, \Delta t_j)$ is not a *signal energy* measure but a measure for the signal bias.

Although Equation (28) minimizes (27), the major issue is that the image irradiance E_i is *unknown* in practice. Therefore we wish to express the weighting function in terms of the observed z_{ij} , and not in terms of the unknown E_i . The pixel intensity z_{ij} (which is a random variable) statistically depends on the image irradiance E_i , through (4). Hence, we can rewrite the weighting function as:

$$\begin{aligned} w_{ij}^{(1)} &= E \left[\frac{A \sigma^2 (E_i, \Delta t_j)}{\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j)} \middle| z_{ij} \right] \\ &= \int_0^{+\infty} \frac{A \sigma^2 (E_i, \Delta t_j)}{\sigma^2 (E_i, \Delta t_j) + P v^2 (E_i, \Delta t_j)} f_{E_i|z} (E_i | z_{ij}) dE_i \end{aligned} \quad (29)$$

where the conditional PDF $f_{E_i|z} (E_i | z_{ij})$ can be computed through Bayes' rule. This requires the specification of two other probability density functions:

1. The prior PDF of the image irradiance $f_E (E_i)$, for which we use a prior distribution of maximal entropy since little information about the image

irradiance is known: the *exponential* distribution. The parameter of the exponential distribution (i.e., the average image irradiance) is assumed to be prior knowledge. To verify whether the exponential distribution is representative for real world images, we performed a simple experiment using three HDR images (in particular, the images from Figures 1d, 11a, and 12a). The obtained histogram and the fitted exponential distribution are shown in Figure 2b. It can be noted that there is a relatively good correspondence (i.e., the distribution well captures the high positive skew). Alternatively, the gamma distribution or even mixtures of gamma distributions may be used as in [25] (note in this respect that the exponential distribution arises as a special case).

2. The conditional PDF $f_{z|E}(z_{ij} | E_i)$, which can be computed exactly from (4) through the change of variables method. More specifically, we have:

$$g(z_{ij} | E_i) \sim N(\log(\sqrt{\alpha_j} \Delta t_j E_i) + v(E_i, \Delta t_j), \sigma^2(E_i, \Delta t_j)),$$

such that the PDF $f_{z|E}(z_{ij} | E_i)$ is given by:

$$f_{z|E}(z_{ij} | E_i) = \varphi(g(z_{ij}), \log(\sqrt{\alpha_j} \Delta t_j E_i) + v(E_i, \Delta t_j), \sigma^2(E_i, \Delta t_j)) \left| \frac{dg}{dz} \right|, \quad (30)$$

with $\varphi(x, \mu, \sigma)$ the Gaussian PDF.

Remark that the weighting function (29) depends on the parameters Δt_j , α_j , β_{ij} , $\sigma_{\sigma_j}^2$ and the CRF $\gamma(x)$. We will explain in Section “Dealing with unknown camera model parameters” how to deal with the scenario in which some of these parameters are also unknown.

Approximate direct formula

The optimal weighting function found in the previous section, depends on the functions $\sigma^2(E_i, \Delta t_j)$ and $v^2(E_i, \Delta t_j)$ and its practical computation requires a (numerical) integration over the image irradiance variable E_i . We therefore investigated if it is possible to find a good approximation to (29), that is somewhat easier to compute.



(a) HDR from densely sampled LDR images.



(b) Debevec (Est. SNR=13.73dB)

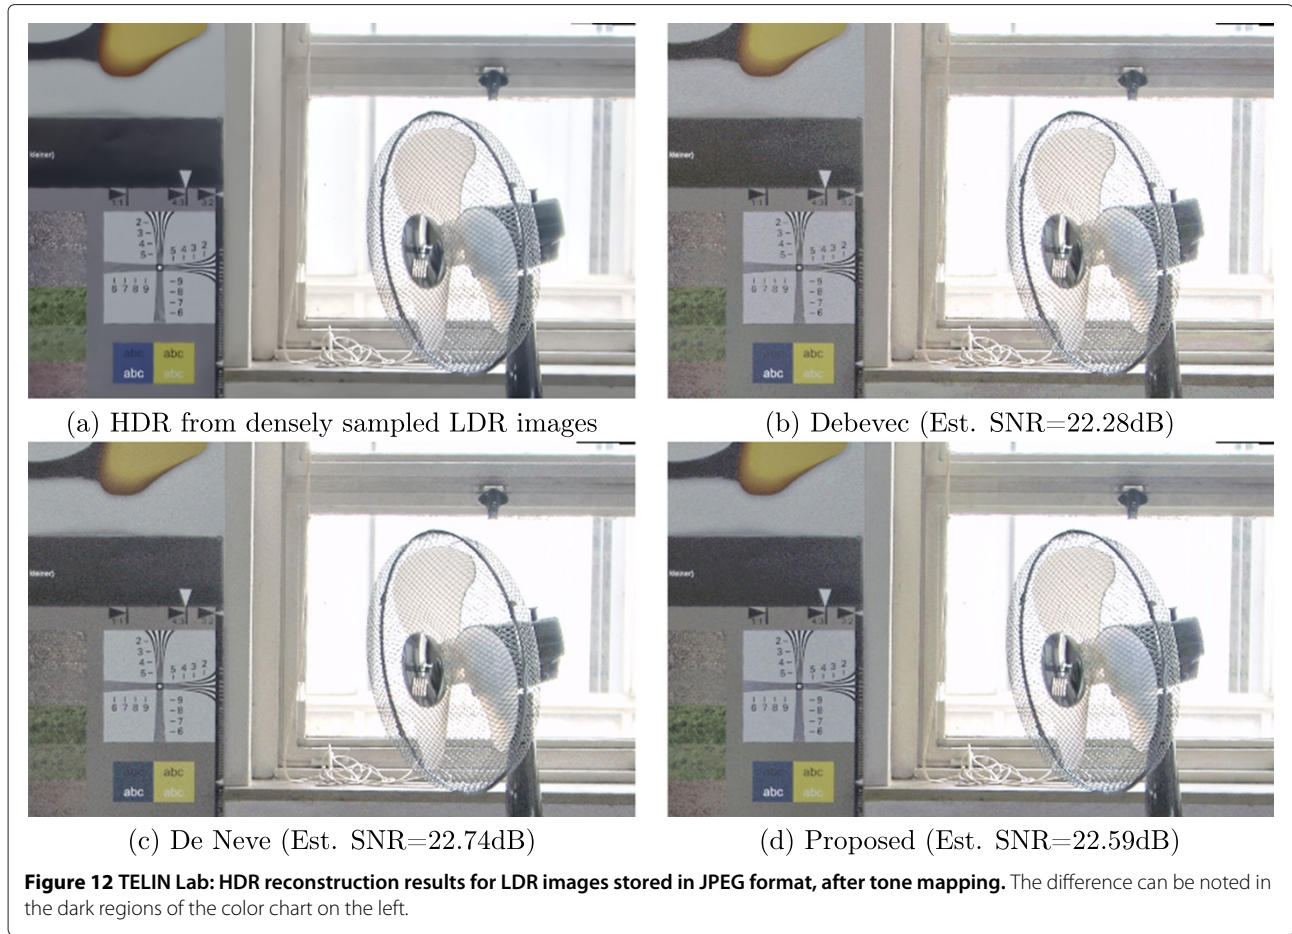


(c) De Neve (Est. SNR=13.71dB)



(d) Proposed (Est. SNR=14.75dB)

Figure 11 Desk still life: HDR reconstruction results for RAW data, after exposure compensation, while balancing and tone mapping.



First, we note that, because the integration range is larger than the region $y'_{\min} \leq \Delta t_j E_i \sqrt{\alpha_j} \leq y'_{\max}$, the approximations (20) and (21) can not be used to simplify (29). However, within the range $y'_{\min} \leq \Delta t_j E_i \sqrt{\alpha_j} \leq y'_{\max}$, we have that $\nu(E_i, \Delta t_j) \approx 0$, such that also $A \approx 1$. Let us denote $E_{\min}(i, j) = y'_{\min}(i, j) / (\Delta t_j \sqrt{\alpha_j})$ and $E_{\max}(i, j) = y'_{\max}(i, j) / (\Delta t_j \sqrt{\alpha_j})$ (assuming $E_{\max} > E_{\min}$), then weighting function (29) can be simplified to:

$$\begin{aligned}
 w_{ij}^{(1)} \approx & \int_0^{E_{\min}(i,j)} \frac{\sigma^2(E_i, \Delta t_j)}{\sigma^2(E_i, \Delta t_j) + P\nu^2(E_i, \Delta t_j)} f_{E|z}(E_i | z_{ij}) dE_i \\
 & + \int_{E_{\max}(i,j)}^{+\infty} \frac{\sigma^2(E_i, \Delta t_j)}{\sigma^2(E_i, \Delta t_j) + P\nu^2(E_i, \Delta t_j)} f_{E|z}(E_i | z_{ij}) dE_i \\
 & + \int_{E_{\min}(i,j)}^{E_{\max}(i,j)} f_{E|z}(E_i | z_{ij}) dE_i \quad (31)
 \end{aligned}$$

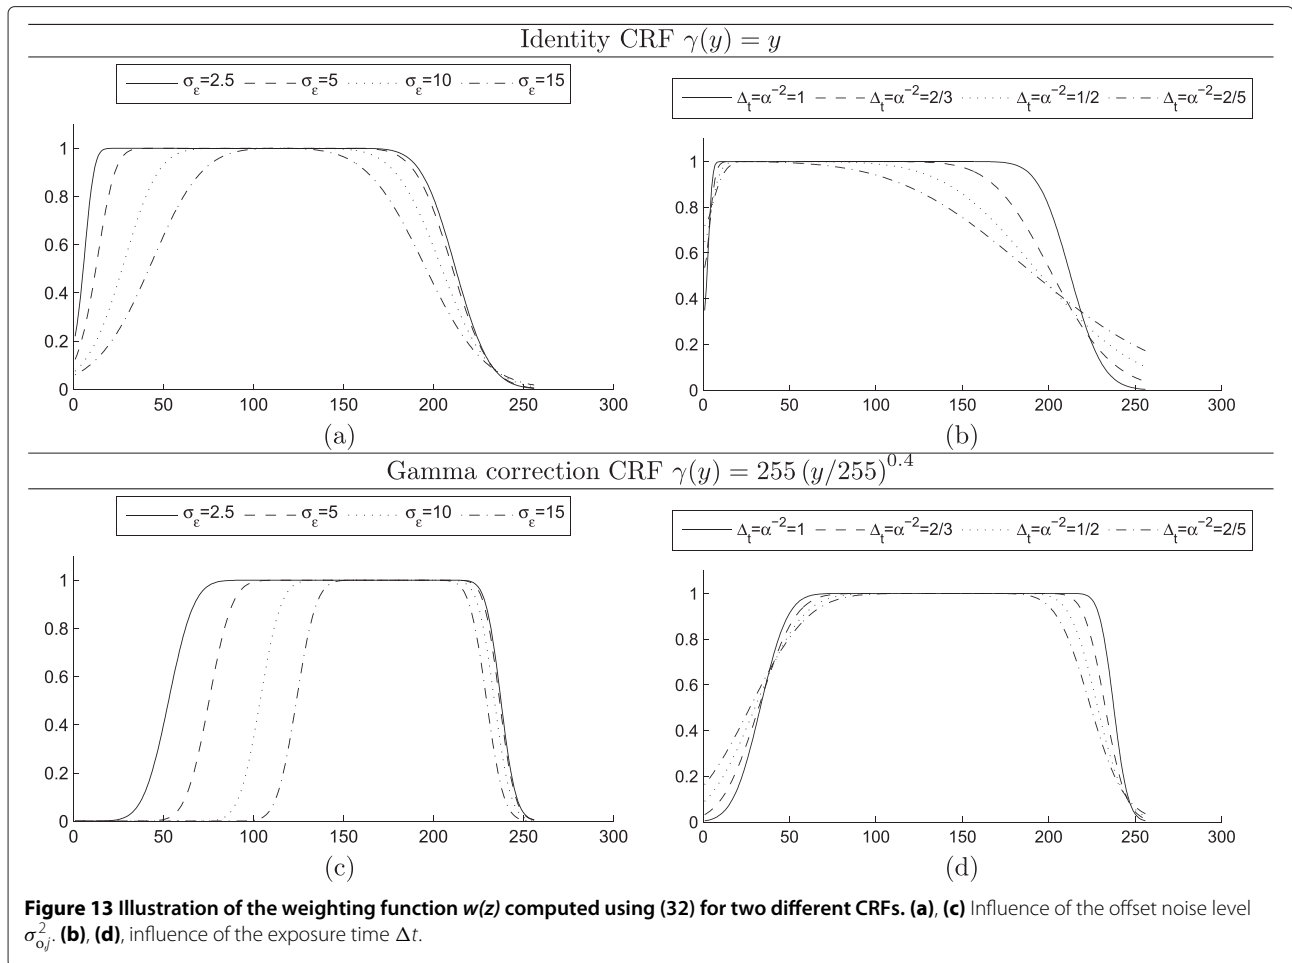
where the approximation error can be controlled by k . We can now define a new weighting function $w_{ij}^{(2)} < w_{ij}^{(1)} / A$ that does not contain any terms in $\sigma^2(E_i, \Delta t_j)$ and

$\nu^2(E_i, \Delta t_j)$ and is hence easier to compute. This leads to our probabilistic formula for the weighting function:

$$\begin{aligned}
 w_{ij}^{(2)} &= \int_{E_{\min}(i,j)}^{E_{\max}(i,j)} f_{E|z}(E_i | z_{ij}) dE_i \\
 &= \Pr[E_{\min}(i, j) < E_i < E_{\max}(i, j) | z_{ij}] \quad (32)
 \end{aligned}$$

which is the probability that the image irradiance is within the bounds $[E_{\min}(i, j), E_{\max}(i, j)]$, given the the observed pixel intensity z_{ij} . Generally, (32) has a higher cost in terms of MSE (because we disregard two terms of (32)), however the squared bias term of (24) will be approximately zero. If $E_i \in [E_{\min}(i, j), E_{\max}(i, j)]$ then the bias $\nu(E_i, \Delta t_j) \approx 0$, so weighting function (32) can be interpreted as “the probability that the inversely compensated pixel intensity $g(z_{ij})$ is approximately bias-free”. It is clear that HDR reconstruction using (32) will also be approximately bias-free: the bias error $E[\log \widehat{E}_i - \log E_i] \approx 0$.

In Figure 13, an illustration of the weighting function is depicted, for two CRFs and for different choices of the parameters. It can be noted that increasing the off-set noise level (resulting in biases with larger magnitudes



$|\nu(E_i, \Delta t_j)|$) inherently decreases the average weight. The same situation occurs for the large intensities (e.g., $z_{ij} > 128$) when decreasing the exposure time. According to the formulas, the asymmetry can be attributed to the Poisson characteristics of the photon noise.

In Appendix 4, a direct formula for (32) was derived:

$$\begin{aligned}
 1w_{ij}^{(3)} &= \frac{\Phi(y'_{\min}(i, j)) - \Phi(y'_{\max}(i, j))}{\Phi(0) - 1} \text{ with } \Phi(x) \\
 &= \operatorname{erf}\left(\frac{x - \gamma^{-1}(z_{ij}) + \lambda\alpha_j^{-1/2}\Delta t_j^{-1}s(\gamma^{-1}(z_{ij}))}{\sqrt{s(\gamma^{-1}(z_{ij}))}}\right).
 \end{aligned}
 \tag{33}$$

The merit of this approximation is that it is somewhat easier to implement and compute in practice. In case we can ignore the quadratic fixed pattern noise (see Section “Camera noise modeling”), Equation (33) only depends on the variables z and j . This allows offline computation and the use of a lookup table.

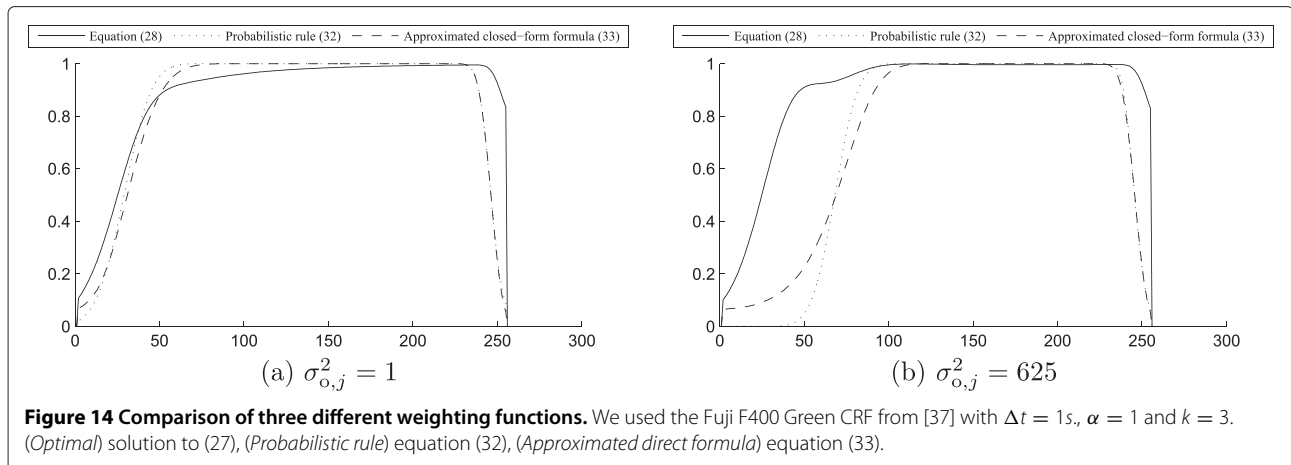
In Figure 14, our three different weighting functions (respectively Equations (28), (32) and (33) are compared

to each other. It can be noted that, generally Equation (28) assigns the largest weight of the three weighting functions. This is because of the approximate relationship $w_{ij}^{(2)} < w_{ij}^{(1)}/A$. Weighting functions (32) and (33) are more *conservative* in assigning large weights to pixel intensities.

Dealing with unknown camera model parameters

In the previous sections, we explained how to determine a parametric weighting function that mediates a trade-off between the SNR of the reconstructed HDR image on the one hand and the squared bias error on the other hand. The function depends on the parameters Δt_j , α_j , β_{ij} , $\sigma_{o,j}^2$ and the CRF $\gamma(x)$, hence beforehand, we assumed that the CRF was known.

In practice, the exposure time can always be assumed to be available.^h The CRF, however, is not publicly disclosed by camera manufacturers, but can be obtained through calibration procedures (e.g., using a color chart), or estimated from a set of LDR images. We will explain this in more detail in Section “Estimation of the CRF”.




If the CRF is available, the camera noise parameters $\alpha_j, \beta_{ij}, \sigma_{o,j}^2$ can be estimated directly using a (robust) linear regression on the the inversely compensated pixel intensities $\gamma^{-1}(z_{ij})$, relying on (17).

One of the main difficulties in estimating the CRF in a noise-robust way is that it is required that the weighting function is known. As we explained, to determine the weighting function, we need the CRF, leading to a chicken-and-egg problem. To resolve this issue, we will now make use of our probabilistic formulation of the weighting function. In Figure 15, it can be seen that range $[y'_{min}, y'_{max}]$ becomes smaller (i.e. y'_{min} increases, y'_{max} decreases) when (1) the offset noise variance $\sigma_{o,j}^2$ increases, (2) the gain factor α decreases or (3) the gain fixed pattern noise factor β_{ij} increases. It can be shown that these observations also hold for $E_{min}(i, j) = y'_{min}(i, j) / (\Delta t_j \sqrt{\alpha_j})$ and $E_{max}(i, j) = y'_{max}(i, j) / (\Delta t_j \sqrt{\alpha_j})$.¹ Since weighting function (32) is computed as the conditional probability that $E_{min}(i, j) < E_i < E_{max}(i, j)$, and because probabilities are always positive, the overall effect is that the weights

$w_{ij}^{(2)}$ decrease under the aforementioned conditions. This behavior of the weighting function can also be noted in Figure 13a,c: curves corresponding to increasing $\sigma_{o,j}^2$ are positioned under each other.

The main idea of our approach is: when little is known about the noise model parameters $\alpha_j, \beta_{ij}, \sigma_{o,j}^2$, we can always select *reasonable upper bounds* (or in case of α_j , an upper bound for the reciprocal α_j^{-1}) for the values of these parameters. The obtained weighting function will then allow some *uncertainty* on the noise model parameters. The only restriction is that the upper bounds have to be chosen such that $E_{max}(i, j) > E_{min}(i, j)$, otherwise all weights become zero (the input SNR would likely be too low to allow proper HDR reconstruction). By using such upper bounds, the weighting function will be more conservative in assigning the weights, but still keeping the bias error (26) small in magnitude. In this sense, the *certainty* function establishes a novel meaning to this concept by relating to the uncertainty associated with the noise model parameters.



Practically speaking, we can fix an upper bound for one parameter, then compute an upper bound for the other parameter values, using the relationship $y'_{\max}(i, j) - y'_{\min}(i, j) = \Delta$ with $\Delta > 0$ a positive number. In Figure 16, we computed the weighting functions according to this parameter uncertainty approach, when varying the lower bound for α for $\Delta = 10$. The gain fixed pattern term β_{ij} was chosen to be 0 here, and an identity CRF $\gamma(x) = x$ was assumed. A simple heuristic using the breaking point from Figure 9, then allows us to determine whether the parameter choice corresponds to the offset noise limited or photon noise limited case. Interestingly, we find that for the offset noise limited case, the weighting function is symmetrical, close to the weighting function of Debevec and Malik and almost equal to the Gaussian weighting function used in [38] for estimating CRFs. Unsurprisingly, the Gaussian weighting function arises as a limit for (33) when Δ is infinitesimally small. On the other hand, in the photon limited case, the weighting function tends to be asymmetrical with a decentralized maximum.

Estimation of the CRF

In this section, we revisit another common problem in HDR reconstruction: the estimation of the CRF. The CRF is often not known in advance, because this requires exact knowledge of the processing steps of the digital cameras and their parameter values, information that is only at hand of camera manufacturers. Therefore it is useful to estimate the CRF in a camera-independent way. In literature, several techniques have been proposed to estimate the CRF. Mann and Picard propose a parametric regression method based on the comparagram (which is a joint histogram of z_{ij} versus $z_{ij'}$ for $j \neq j'$) [27,30]. Mitsunaga and Nayar [29] perform an iterative polynomial regression, where the (assumed to be unknown) exposure time ratio $\Delta t_j / \Delta t_{j'}$ is refined in each iteration. Related iterative methods have been proposed by Tsin et al. [33]. Lin et al. [39] use a completely different technique that estimates

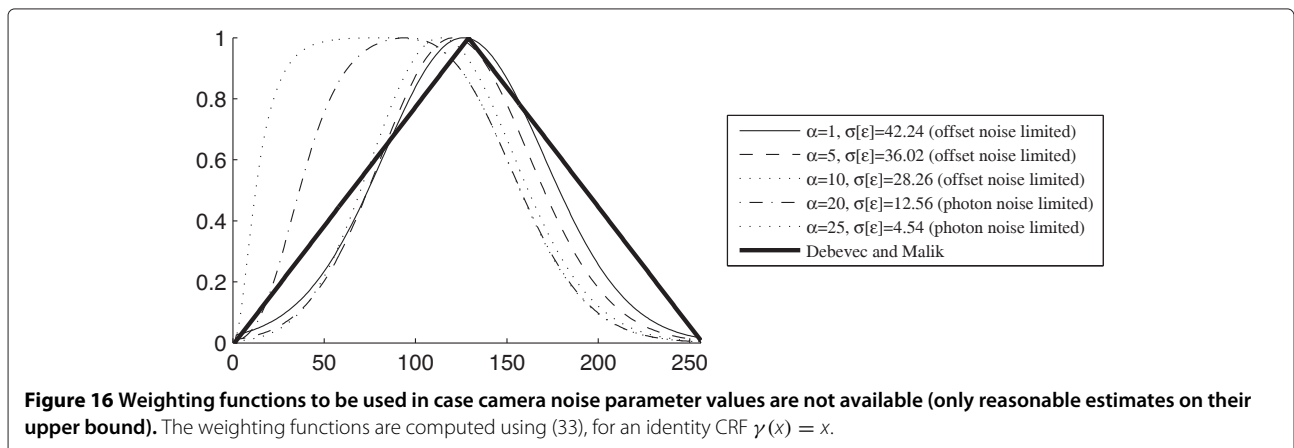
the CRF based on RGB distributions of pixels at color edges, requiring only one exposure ($P = 1$). Debevec and Malik use a non-parametric method to estimate the inverse CRF $g(z)$ [28], leaving a lot of degrees of freedom to this function. In [31], this method was improved to incorporate all available pixel intensities in the estimation (leading to more accurate estimation), while reducing computational complexity.

The problems with the existing techniques are that they are either not very robust to high levels of noise (for example, the method of Debevec and Malik does not enforce monotonicity in terms of a direct constraint to the CRF estimation problem, often leading to CRF estimates that are oscillating in presence of noise) or not very well adapted to the signal-dependent noise characteristics of the individual LDR images (e.g., due to an assumption of stationary noise).

Let us analyze the CRF estimation problem by relying on the camera model developed in the previous sections. Consider two distinct LDR images z_{ij} and $z_{ij'}$. When z_{ij} and $z_{ij'}$ are in the appropriate exposure range ($y'_{\min}(i, j) < \alpha_j \Delta t_j E_i < y'_{\max}(i, j)$ and $y'_{\min}(i, j') < \alpha_{j'} \Delta t_{j'} E_i < y'_{\max}(i, j')$), we have $E[z_{ij}|E_i] \approx \gamma(\sqrt{\alpha_j} \Delta t_j E_i)$ and $E[z_{ij'}|E_i] \approx \gamma(\sqrt{\alpha_{j'}} \Delta t_{j'} E_i)$ with respective variances (see Appendix 1):

$$\begin{aligned} \text{Var}[z_{ij}|E_i] &\approx \left(\sigma_{\alpha_j}^2 + \alpha_j \Delta t_j E_i + \beta_{ij} (\Delta t_j E_i)^2 \right) \left(\frac{\partial \gamma}{\partial x} \right)^2 \Bigg|_{x=\Delta t_j E_i} \\ \text{Var}[z_{ij'}|E_i] &\approx \left(\sigma_{\alpha_{j'}}^2 + \alpha_{j'} \Delta t_{j'} E_i + \beta_{ij'} (\Delta t_{j'} E_i)^2 \right) \left(\frac{\partial \gamma}{\partial x} \right)^2 \Bigg|_{x=\Delta t_{j'} E_i} \end{aligned} \quad (34)$$

When considering the estimation of the CRF as a regression problem (as in [30]), optimal estimation is



quite difficult due to the dependency of the variance on the (*unknown*) image irradiance E_i . If the conditional variances $\text{Var}[z_{ij}|E_i]$ and $\text{Var}[z_{ij'}|E_i]$ are constant (e.g. $\text{Var}[z_{ij}|E_i] \approx \sigma_{\alpha_j}^2$), the method of total least squares (TLS) [40] is well suited to estimate the nonlinear CRF: the TLS method performs a regression in which the *orthogonal* distance error of the fitted function to the observations is minimized. The TLS method is well suited for regression problems in which both the x-variables and y-variables have a known and constant variance. Unfortunately, this is not the case here. Furthermore, the variances (34) depend on the derivative of the CRF, which is to be estimated!

As an alternative solution, we can rely on the statistical conditional independence of z_{ij} and $z_{ij'}$ (given E_i). According to (4), the difference $g(z_{ij}) - g(z_{ij'})$ has the following PDF:

$$g(z_{ij}) - g(z_{ij'}) \sim \mathcal{N} \left(\log \frac{\alpha_j \Delta t_j}{\alpha_{j'} \Delta t_{j'}} + v(E_i, \Delta t_j) - v(E_i, \Delta t_{j'}), \sigma^2(E_i, \Delta t_j) + \sigma^2(E_i, \Delta t_{j'}) \right).$$

In the exposure range $\{\alpha_j \Delta t_j E_i, \alpha_{j'} \Delta t_{j'} E_i\} \subset [y'_{\min}(i, j), y'_{\max}(i, j')]$, the bias terms can safely be ignored (provided that E_i is sufficiently large):

$$g(z_{ij}) - g(z_{ij'}) \sim \mathcal{N} \left(\log \frac{\Delta t_j}{\Delta t_{j'}} + \log \frac{\alpha_j}{\alpha_{j'}}, \sigma^2(E_i, \Delta t_j) + \sigma^2(E_i, \Delta t_{j'}) \right). \quad (35)$$

Maximum likelihood estimation of $g(z_{ij})$ now amounts to the linear regression:

$$\widehat{g}(z) = \min \sum_{j'=1}^P \sum_{i=1}^N w_{ij} w_{ij'} \frac{|g(z_{ij}) - g(z_{ij'}) - \log \frac{\Delta t_j}{\Delta t_{j'}} - \log \frac{\alpha_j}{\alpha_{j'}}|^2}{\sigma^2(E_i, \Delta t_j) + \sigma^2(E_i, \Delta t_{j'})}, \quad (36)$$

where a weighting function was introduced in analogy to Equation (6). Recall that the LDR pixel intensities are discrete ($z_{ij} \in 0, \dots, 2^B - 1$ and $z_{ij'} \in 0, \dots, 2^B - 1$). Therefore, (36) can be solved by treating $g(z_{ij})$ and $g(z_{ij'})$ as unknown variables. Solving (36) then gives a linear system of $N \times P$ equations with 2^B unknowns ($g(0), \dots, g(2^B - 1)$).

In case either offset noise or photon noise is dominant, the sum $\sigma^2(E_i, \Delta t_j) + \sigma^2(E_i, \Delta t_{j'})$ can be approximated by (see (22)):

$$\sigma^2(E_i, \Delta t_j) + \sigma^2(E_i, \Delta t_{j'}) \approx \begin{cases} \frac{1}{E_i} \frac{\Delta t_j + \Delta t_{j'}}{\Delta t_j \Delta t_{j'}} & \text{(photon noise dominant)} \\ \frac{\sigma_{\alpha_j}^2 \Delta t_j^2 \alpha_j + \Delta t_{j'}^2 \alpha_{j'}}{E_i^2 \alpha_j \alpha_{j'} \Delta t_j \Delta t_{j'}} & \text{(offset noise dominant)} \end{cases}.$$

As we mentioned earlier, because of estimation in presence of noise, we need to explicitly enforce that $g(z)$ is monotonic. Therefore, we include the conditions $g(z) > g(z - 1)$, $z = 1, \dots, 2^B - 1$ as extra constraints to the minimization problem from (36). Reorganizing the terms in (36) then gives (in case photon noise is dominant in all of the LDR images):

$$\min \sum_{i=1}^N \frac{1}{E_i} \sum_{j'=1}^P w_{ij} w_{ij'} \frac{|g(z_{ij}) - g(z_{ij'}) - \log \frac{\Delta t_j}{\Delta t_{j'}} - \log \frac{\alpha_j}{\alpha_{j'}}|^2}{(\Delta t_j + \Delta t_{j'}) / \Delta t_j \Delta t_{j'}} + \lambda \sum_{z=0}^{2^B-1} w_{\text{smooth}}^2(z) \left| \frac{d^2 g}{dz^2} \right|$$

s.t. $g(z) \geq g(z - 1) + \epsilon$, $z = 1, \dots, 2^B - 1$. (37)

with λ a regularization parameter to enforce smoothness of $g(z)$ (see [28,31]), and with ϵ a small positive number. In (37), the function $w_{\text{smooth}}(z)$ is the weighting function, averaged over the set of images $j = 1, \dots, P$. To compute $d^2 g / dz^2$, the numerical second derivative is used. Optimization problem (37) can be efficiently solved using standard quadratic program (QP) solvers [41]. We note that in (37), the (unknown) image irradiance appears as position-dependent weights. A straightforward solution is then to iteratively update the CRF and the log-irradiance estimate $\log E_i$, but to keep the processing technique simple, we propose here to simply drop the position-dependent weights $1/E_i$. To conclude, the proposed CRF estimation technique differs from the technique in [31] in the following aspects: (1) an image-dependent weight factor $(\Delta t_j + \Delta t_{j'}) / \Delta t_j \Delta t_{j'}$ that penalizes LDR images with low exposure times (because the SNR is generally lower in these images), (2) the use of a weight function w_{ij} that is adapted to the underlying camera noise model and (3) the monotonicity constraint for $g(z)$. Consequently, the CRF estimation technique will be considerably more robust against noise, with only limited increase in algorithmic complexity (the use of a QP solver compared to a sparse system solver).

Results and discussion

Comparison of different weighting functions

In this section, we compare the different weighting functions proposed in literature to our probabilistic function from Equation (32). We will only consider the influence of the choice of the weighting function on the HDR reconstruction (i.e., we do not include the CRF estimation in this experiment and all techniques use the same HDR reconstruction formula from (7). As in Equation (7), the effective weights are calculated using $w_{ij} \sigma^{-2}(E_i, \Delta t_j)$. We computed the performance of the different weighting functions, using analytical formulas (Equations (24)–(26), for different camera gains $\sqrt{\alpha} \in \{0.25, 0.5, 1, 2, 4\}$ and different offset noise levels $\sigma_o \in$

{0.001, 0.5, 1, 2, 4, 8, 16}. We assumed four LDR images with exposure times 0.1s, 0.25s, 1s, 2s and an identity CRF $\gamma(z) = z$. The expected MSE $E \left[\left(\widehat{\log E} - \log E \right)^2 \right]$ is inherently averaged over a logarithmically spaced range over the image irradiance E , covering the whole input dynamic range. We remark that an upper bound for the theoretical maximal gain compared to the triangular weighting function of Debevec and Malik is given by:

$$10 \log_{10} \frac{\int_0^1 z^2 dz}{\left(\int_0^1 z dz \right)^2} = 10 \log_{10} \frac{4}{3} = 1.25 \text{ dB.}$$

This means that, for any estimation task, no weighting function will give a SNR improvement of more than 1.25 dB, compared to the triangular weighting function. The above formula ignores the clipping of the dynamic range and the presence of signal-dependent noise, hence in practice, the actual improvement will be much smaller than 1.25 dB. The actual values are plotted in Figure 17, it can be seen that the probabilistic weighting function reaches a maximal improvement of approximately 1 dB (for larger values of σ_o^2) and outperforms the other weighting functions in most of the cases.

Ground truth data with simulated noise

To have ground truth data, we generated a computer rendering of a HDR scene (see Figure 18a). The scene consists of a room with a textured walls with four paintings

(all textures and paintings have a bit depth of 8). In the middle of the room, a bright point light source is positioned, creating a dynamic range of 11 stops (i.e. the ratio of the brightest and smallest possible pixel intensity is about 2^{11}). Sensor signals were simulated and artificial noise was thereby generated, according to the procedure explained in Section “Camera noise modeling”. An example of such exposes LDR images is given in Figure 18b–e. The light source causes parts of the wall, floor and ceiling to be overexposed in some of the LDR images, while the wall on the right is underexposed in other images. Results are again generated for different camera gains $\sqrt{\alpha} \in \{0.25, 0.5, 1, 2, 4\}$ and different offset noise levels $\sigma_o \in \{0.001, 0.5, 1, 2, 4, 8, 16\}$, for a CRF modeling gamma correction $\gamma(z) = z^{0.4}$ under (1) the assumption that the CRF is known *in advance*, and (2) by estimating the CRF from the images. For the method of Debevec and Malik [28], the CRF was estimated from 1,000 samples, randomly sampled in the image. The regularization parameter is fixed to $\lambda = 1$ for all methods. The results are given in Figure 19. It can be seen that, in case the CRF is known, an improvement of about 1dB is obtained by adapting to the camera noise model. When also CRF estimation is included in the algorithm, the SNR further improves by 1–1.5 dB. Also notice that the weighting function from Section “Dealing with unknown camera model parameters” that includes uncertainty with respect to the noise model parameters gives slightly worse results than the proposed probabilistic weighting function.

A visual HDR reconstruction result is presented in Figure 20. In the image reconstructed using the proposed technique, the noise is significantly better suppressed.

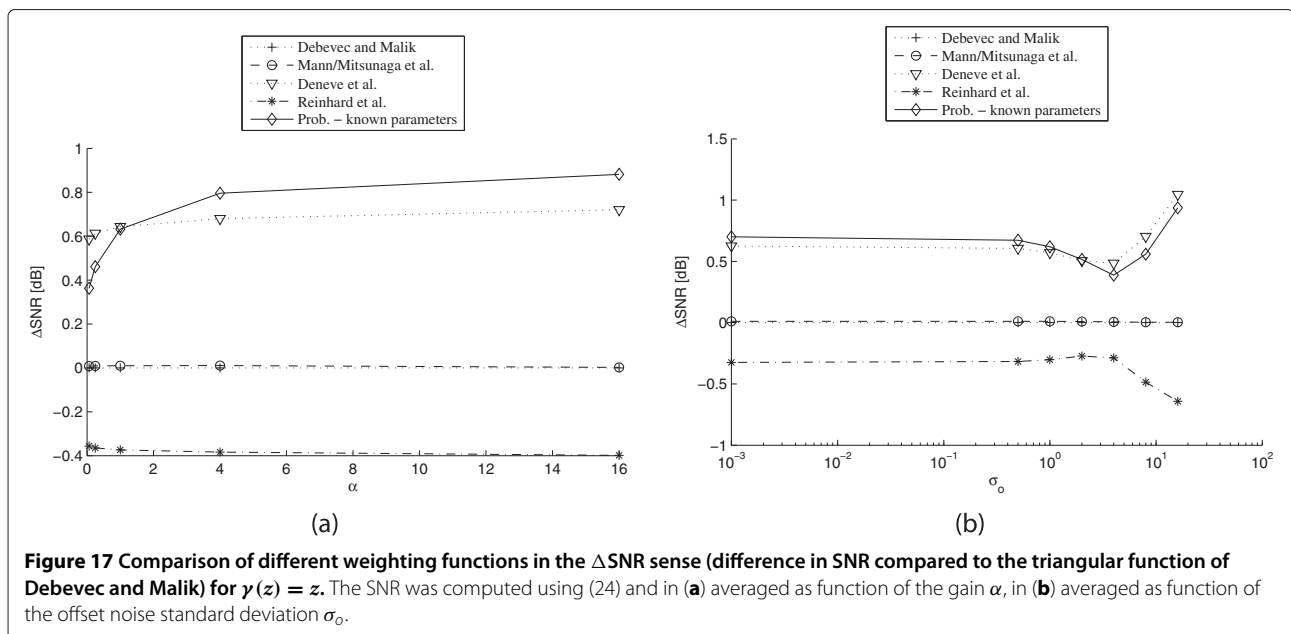


Figure 17 Comparison of different weighting functions in the Δ SNR sense (difference in SNR compared to the triangular function of Debevec and Malik for $\gamma(z) = z$. The SNR was computed using (24) and in (a) averaged as function of the gain α , in (b) averaged as function of the offset noise standard deviation σ_o .

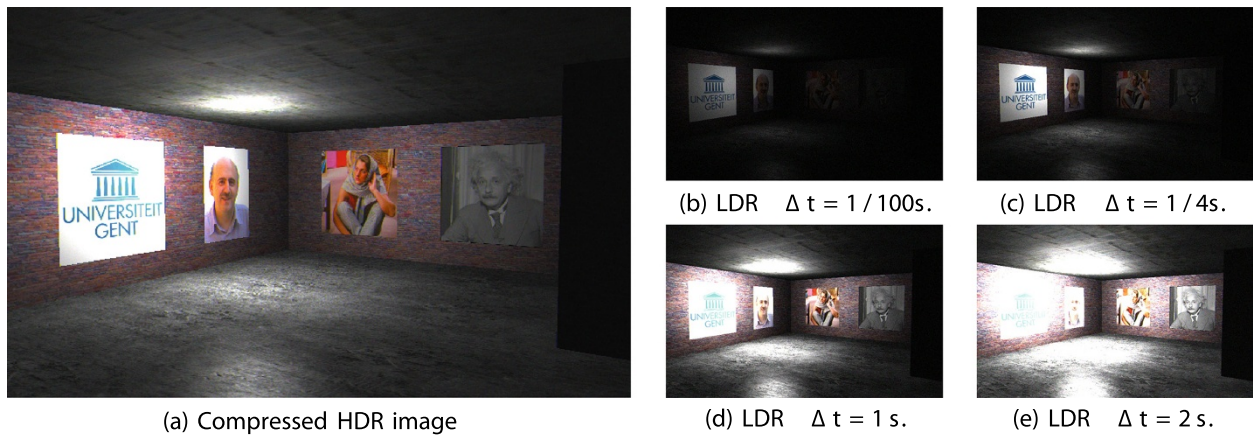


Figure 18 Ground truth data used for the experiments. **(a)** HDR image with compression from [26]. **(b)–(e)** LDR images generated from **(a)** using different exposure times.

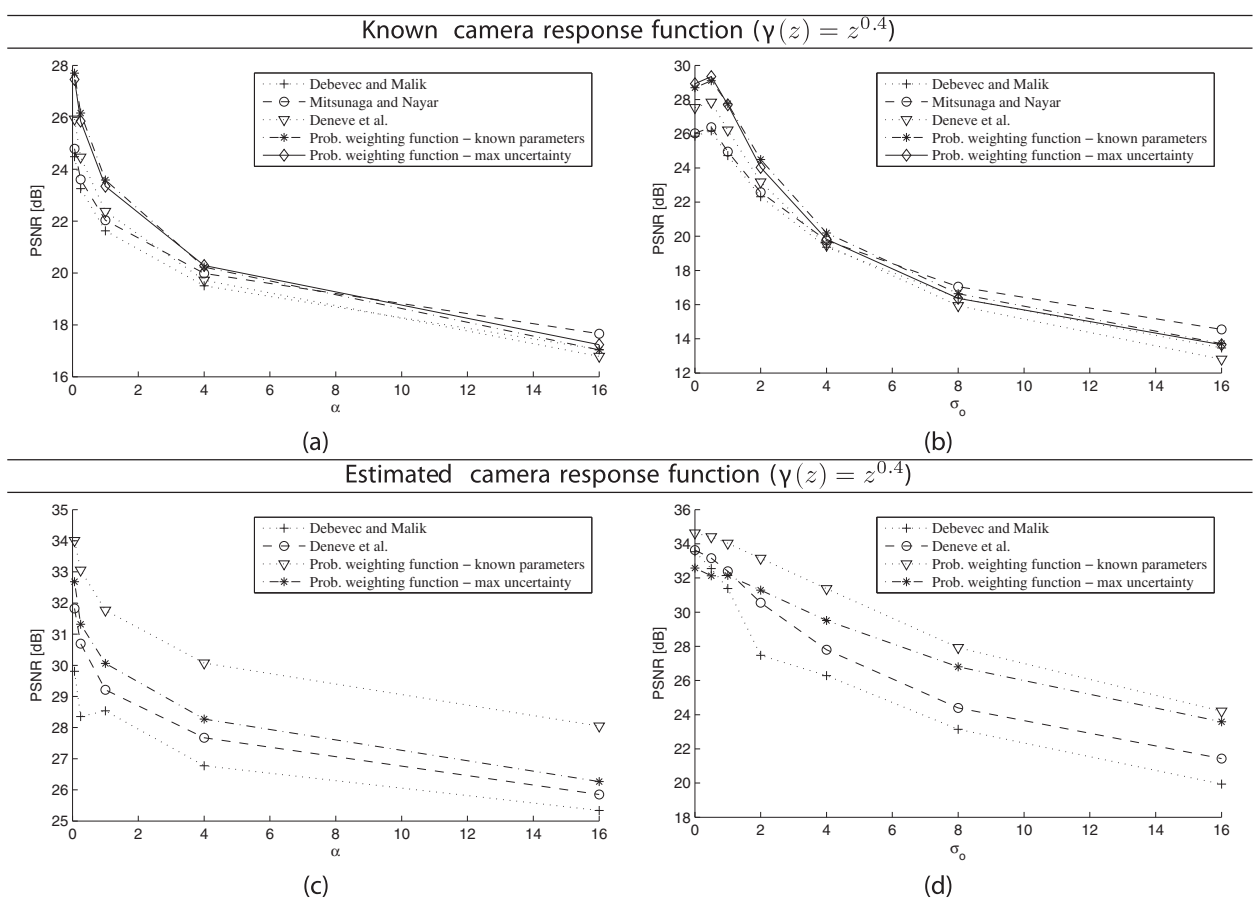


Figure 19 HDR reconstruction results using ground truth data. **(a)** Average PSNR results as function of gain α , **(b)** average PSNR results as function of the offset noise standard deviation σ_o .



(a) Weighting function of Debevec et al.: PSNR = 25.70dB. (b) Prob. weighting function PSNR = 28.59dB.

Figure 20 Visual comparison of HDR reconstruction using different weighting functions ($\alpha = 1/64$, $\beta = 0$, $\sigma_{\omega_j}^2 = 64$, $j = 1, \dots, 4$) and different CRF estimation methods.

Although the maximal SNR improvement of the HDR image is bounded when only using position-dependent averaging (see Section “Denoising performance of the HDR reconstruction formula”), further improvement can be obtained by including spatial regularization in the reconstruction process.

HDR reconstruction of RAW sensor data

Next, we demonstrate our method for real digital camera images. In particular, we use the “desk still life” set of 17 RAW LDR images acquired by Hasinoff [24] using a Canon EOS 1D Mark III camera (10 mega-pixel images, with an equal ISO setting of 100).^j The 17 images are used to obtain an estimate for the ground-truth data, which is useful for objective evaluation. The spatial resolution of the images and the exposure time sampling is sufficiently high, such that different parameters of the reconstruction methods have a limited impact on the reconstruction result. Here, we use Debevec’s method[28] for creating the HDR images.

To compare different methods, we select 3 LDR images out of the 17 images, with exposure times 1/800, 1/200,

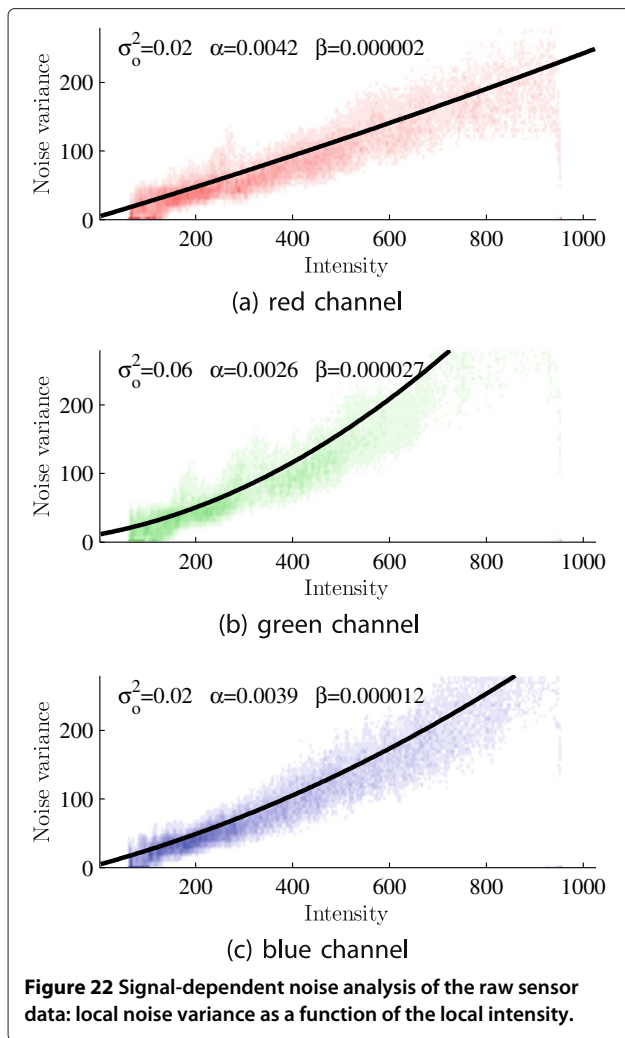
1/25 s. These exposure times are chosen to cover a large portion of the dynamic range of the scene, such that some of the LDR images exhibit high levels of noise (especially in the dark regions, see Figure 21). Although this choice is somewhat arbitrary, similar results can be obtained using slightly deviating exposure time values.

First, to estimate the camera noise model parameters (see Section “Camera noise modeling”), we analyze signal-dependent noise in a similar way as in [42], i.e., by computing the local noise standard deviation as a function of the local intensity. We use a 7×7 local window to estimate the local mean and standard deviation, where we also exclude patches in which the signal is too strong (to find these patches, we compare the maximum local magnitude of the Laplacian filtered image to a fixed threshold, and we reject patches if the maximum is too large). We repeat this for the three color channels, and we directly fit the parabolic noise variance model (17) to the obtained point clouds using least squares fitting. The noise fitting results are shown in Figure 22 for the LDR image of exposure time 1/25 s. It can be seen that there is a good correspondence between the predicted noise variance and the



(a) $\Delta t_1 = 1/800s$ (b) $\Delta t_2 = 1/200s$ (c) $\Delta t_3 = 1/25s$

Figure 21 Desk still life: LDR images (after camera reconstruction).



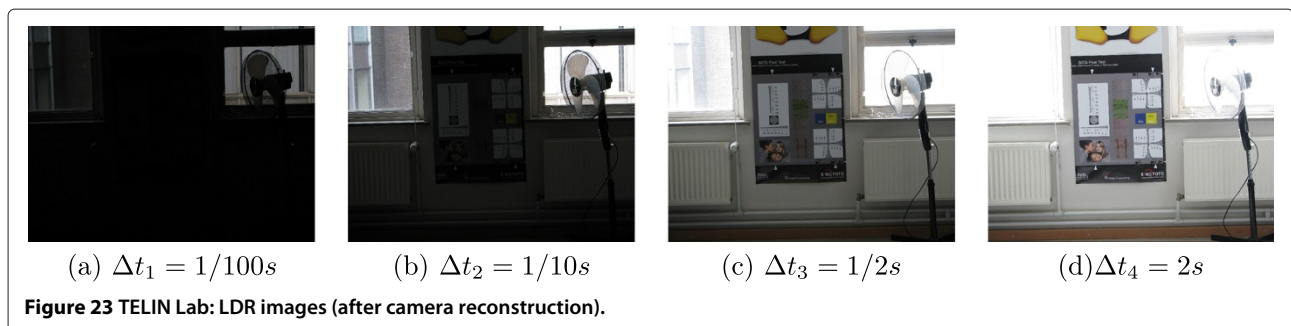
estimated noise variance. Alternatively, affine fitting could also be used here (assuming the absence of gain FPN, $\beta_{ij} = 0$ in (17), however, we found that in most LDR images a quadratic model better describes the measured noise variances.

A HDR image is then constructed from the three LDR images, using the proposed approach or, alternatively,

existing methods like [28,31]. Finally, some common post-processing operations, such as white balancing, exposure correction, color correction and HDR tone mapping need to be applied to the constructed HDR images. For this post-processing stage, we employ the software program “Raw Therapee V4.0”^k For our purposes, this program allows us to create a reconstruction profile with fixed (non automatic) parameter settings, which can then subsequently be applied to the images reconstructed using different methods. This is very useful to ease visual comparison. In Figure 11, the final results are shown, together with the HDR reconstructed from densely sampled LDR images. The signal-to-noise ratio is estimated from the log-irradiance values obtained before post-processing, by using this last image as a reference. It can be seen that by using optimized weighting functions we can obtain a significant increase in visual quality, compared to other methods which are using more “general-purpose” weighting functions.

HDR reconstruction from JPEG compressed LDR digital photographs

As a last experiment, we test our method on a set of 15 LDR images, with increasing exposure times and a fixed ISO setting of 100 and aperture of f3.5, captured using a Canon Powershot S5 camera, with post-processing and JPEG compression performed by the camera. As in the previous experiment, we use all 15 images to obtain an estimate for the ground-truth data, and we select a subset of four images (with exposure times 1/100, 1/10, 1/2 and 2 s, see Figure 23), for comparing the different HDR reconstruction methods. Because camera post-processing is applied, all methods now require CRF estimation. One of the issues arising for our method, is that the noise model parameter estimation (as used in the previous experiment) is not directly applicable. Note in this respect that applying the inverse CRF is *not* sufficient to obtain estimates of the RAW sensor data, because the component-wise CRF does not take color space operations (such as color corrections, white balancing,...) into account. Instead, we compute the weighting functions as explained in Section



“Dealing with unknown camera model parameters” by defining upper bounds for α_j^{-1} and $\sigma_{o,j}^2$. Here we use the values $\max(\alpha_j^{-1}) = 0.005$, $\max(\sigma_{o,j}^2) = 0.02$ and $\max(\beta_{ij}) = 0$, which are experimentally determined. In our method, the weighting functions used for CRF estimation are based on an identity CRF (recall the chicken-and-egg problem mentioned in Section “Dealing with unknown camera model parameters”). Next, to reconstruct the HDR image, we re-compute the weighting functions using the estimated CRF using results of Appendices 1 and 2. Finally, tone mapping with fixed parameters is applied using the program “Raw Therapee V4.0”. Some visual results are depicted in Figure 12 and also estimated SNR values (obtained by comparing to the “estimated” ground-truth data) are reported. It can be seen that only small differences exist in the reconstruction results and that the proposed method performs within the same range as the other methods (in particular, the difference can be spotted in the dark regions of the color chart in Figure 12: while Debevec’s method still leaves a lot of noise in this region, the method of [31] and the proposed method do a better job in suppressing the noise). The reason that the proposed method does not outperform the other methods in the same way as in Section “HDR reconstruction from JPEG compressed LDR digital photographs” is most likely a combined effect of several factors: (1) an estimate of the CRF is required to determine the weighting functions for (32), (2) CRF estimation errors also influence the weighting functions and (3) nonlinear color space operations are not taken into account. Although further improvements are possible, e.g., by iteratively refining the weighting functions and CRFs, this experiment already shows that the proposed approach is robust enough to work with JPEG compressed images where standard post-processing operations have been applied.

Conclusion

In this article, we presented a realistic digital camera noise model that incorporates several noise sources (such as photon noise, electronic noise, fixed pattern noise), as well as several parts of the camera processing pipeline (such as the amplifier, clipping, ...). For this noise model, we derived both exact and approximate formulas for the bias function and the noise variance function, taking clipping of the dynamic range into account. Next, we investigated novel HDR image reconstruction weighting schemes relying on the realistic noise model. We paid special attention to different factors that determine the signal-to-noise ratio of the HDR image, and we used the obtained insight to derive

weighting functions that are optimized for the camera noise model. This led to our probabilistic weighting function, that is defined as the probability that the considered pixel intensity is (approximately) unbiased. Because due to the clipping there is a strong coupling between the noise variance and bias functions, this weighting function offers a trade-off between maximizing the SNR of the HDR image and bias errors, its performance is close to optimal in the MSE sense. The probabilistic reasoning was then extended to derive a weighting function for the case in which (some of) the noise model parameters are unknown. Experimental results confirmed the expected improvements in reconstruction quality, especially for reconstruction from RAW sensor data. Although HDR image reconstruction is only one example of application of sophisticated noise modeling techniques, there is a wide range of other application areas to explore in which similar image quality improvements can be obtained when adapting to the noise characteristics from the imaging devices.

Appendix

Appendix 1: approximation of the statistical moments of a non-linear function of signal and noise

In this section, we develop a technique that gives simple expressions that relate the noise statistical moments to the unknown parameters. First, we start from a generative model:

$$u = \zeta(E, v) \quad (38)$$

where E is the image irradiance and v is Gaussian noise with zero mean and unit variance $N(0, 1)$. The signal-and-noise mixing function $\zeta(E, v)$ has the property that it is infinitely differentiable near the working point $v = 0$, for all E . In our application, the mixing function is given by:

$$\zeta(E, v) = \gamma \left(\sqrt{\alpha} \Delta t E + \sqrt{s(\Delta t E)} v \right) \quad (39)$$

with $s(x) = \sigma_{o,j}^2 + \alpha x + \beta x^2$. The conditional statistical moments of u can be computed through the McLaurin series expansion of $\zeta(E, v)$. For the first order moment we have:

$$\begin{aligned} E[u|E] &\approx \zeta(E, 0) + \sum_{n=1}^{+\infty} \left. \frac{\partial^n \zeta}{\partial v^n} \right|_{v=0} \frac{v^n}{n!} \\ &= \zeta(E, 0) + \sum_{n=1}^{+\infty} \left. \frac{\partial^{2n} \zeta}{\partial v^{2n}} \right|_{v=0} \frac{2^{-n}}{n!} \end{aligned} \quad (40)$$

where we used the statistical moments of the Gaussian distribution. A similar expression can be given for the second order moment:

$$\begin{aligned} E[u^2|E] &\approx \zeta^2(E, 0) + \sum_{n=1}^{+\infty} \frac{\partial^n \zeta^2}{\partial v^n} \Big|_{w=0} \frac{E[v^n]}{n!} \\ &= \zeta^2(E, 0) + \sum_{n=1}^{+\infty} \frac{\partial^{2n} \zeta^2}{\partial v^{2n}} \Big|_{w=0} \frac{2^{-n}}{n!} \\ &= \zeta^2(E, 0) + 2 \sum_{n=1}^{+\infty} \frac{2^{-n}}{n!} \\ &\left(\zeta(E, 0) \left(\frac{\partial^{2n} \zeta}{\partial v^{2n}} \Big|_{v=0} \right) + \frac{\partial^{2n-1} \zeta}{\partial v^{2n-1}} \Big|_{v=0} \left(\frac{\partial \zeta}{\partial v} \Big|_{v=0} \right) \right) \end{aligned}$$

Note that $\text{Var}[u|E] = E[u^2|E] - (E[u|E])^2$. The reason for such an expansion is the fact that in practical circumstances, $\zeta(E, v)$ is a relatively smooth function, i.e., the derivatives $\partial^n \zeta / \partial v^n$ evaluated at $v = 0$ decay exponentially as a function of n . For this reason, only one term of the series expansion is often sufficient in our application (provided that $\zeta(E, v)$ is differentiable near $w = 0$). When only considering one term of the McLaurin series approximation, we find:

$$\begin{aligned} E[u|E] &\approx \zeta(E, 0) + \frac{1}{2} \frac{\partial^2 \zeta}{\partial v^2} \Big|_{v=0} \\ \text{Var}[u|E] &\approx \frac{\partial \zeta}{\partial v} \Big|_{v=0}^2 - \frac{1}{4} \frac{\partial^2 \zeta}{\partial v^2} \Big|_{v=0}^2 \approx \frac{\partial \zeta}{\partial v} \Big|_{v=0}^2 \end{aligned}$$

where we used the fact that $\frac{\partial^2 \zeta}{\partial v^2} \Big|_{v=0}^2 \ll \frac{\partial \zeta}{\partial v} \Big|_{v=0}^2$. Substitution of the mixing function (39) gives explicit expressions for the conditional statistical moments of u :

$$E[u|E] \approx \gamma(\sqrt{\alpha} \Delta t E) + \frac{s(\Delta t E)}{2} \frac{\partial^2 \gamma}{\partial x^2} \Big|_{x=\sqrt{\alpha} \Delta t E} \quad (41)$$

$$\text{Var}[u|E] \approx s(\Delta t E) \left(\frac{\partial \gamma}{\partial x} \Big|_{x=\sqrt{\alpha} \Delta t E} \right)^2 \quad (42)$$

where $\gamma(x) = \zeta(x, 0)$ is the CRF. For the special case of a logarithmic transform $\gamma(x) = \log x$, we get:

$$E[u|E] \approx \log(\sqrt{\alpha} \Delta t E) - \frac{1}{2} \frac{s(\Delta t E)}{\alpha \Delta t^2 E^2} \quad (43)$$

$$\text{Var}[u|E] \approx \frac{s(\Delta t E)}{\alpha \Delta t^2 E^2} \quad (44)$$

These convenient formulas form the basis of our noise model in Section ‘‘Camera noise modeling’’.

Appendix 2: determining the clipping-safe region of the dynamic range

In this section, we will determine a region of the dynamic range where the approximations (23) and (21) hold, with high probability. We remark that these approximations are obtained using the technique presented in Appendix 1 and that this technique fails for working points of the clipping function, where this function is not differentiable. Obviously, this is near the saturation points of the dynamic range. We will analytically determine the effect of the clipping on the statistical moments of $\log f(y_{ij})$. First, we define the moment function as:

$$\begin{aligned} M_E(n) &= E[(\log f(y_{ij}))^n | E] \\ &= \frac{1}{\sqrt{2\pi s(\Delta t E)}} \int_{-\infty}^{+\infty} (\log f(u))^n \exp \\ &\left(-\frac{1}{2} \frac{(\sqrt{\alpha} \Delta t E - u)^2}{s(\Delta t E)} \right) du. \end{aligned} \quad (45)$$

The bias and variance functions from (19) can be expressed in terms of $M_E(n)$:

$$\begin{aligned} v(E, \Delta t) &= M_1(E) - \log f(\sqrt{\alpha} \Delta t E) \\ \sigma^2(E, \Delta t) &= M_2(E) - (M_E(1))^2 \end{aligned}$$

Suppose that the clipping function, in a general form, is defined by:

$$f(y) = \begin{cases} y_{\min} & y < y_{\min} \\ y & y_{\min} \leq y \leq y_{\max} \\ y_{\max} & y > y_{\max} \end{cases}$$

we can rewrite the moment function (45) as:

$$\begin{aligned} M_n(E) &= \frac{1}{\sqrt{2\pi s(\Delta t E)}} \int_{-\infty}^{+\infty} (\log u)^n \exp \left(-\frac{1}{2} \frac{(\sqrt{\alpha} \Delta t E - u)^2}{s(\Delta t E)} \right) du \\ &\quad - \frac{1}{2} (\log y_{\min})^n \left(1 + \text{erf} \left(\frac{y_{\min} - \sqrt{\alpha} \Delta t E}{\sqrt{2s(\Delta t E)}} \right) \right) \\ &\quad - \frac{1}{2} (\log y_{\max})^n \left(1 + \text{erf} \left(\frac{-y_{\max} + \sqrt{\alpha} \Delta t E}{\sqrt{2s(\Delta t E)}} \right) \right). \end{aligned} \quad (46)$$

Because the first term of $M_E(n)$ is not related to the clipping function, we can use the technique from Appendix 1 to approximate this term (this term represents the statistical moments as if a clipping function was not present). Because of the decay of $1 + \text{erf}(-x)$, the second and third terms in (46) vanish in the following region of the dynamic range:

$$y_{\min} + k\sqrt{s(\Delta t E)} < \sqrt{\alpha} \Delta t E < y_{\max} - k\sqrt{s(\Delta t E)} \quad (47)$$

with k a positive constant. Next, we can exploit the fact that $s(x) = \sigma_{\alpha_j}^2 + \alpha x + \beta x^2$. To proceed, we use the linear Taylor approximations:

$$\sqrt{s(\Delta t E)} \approx \begin{cases} \sqrt{s(y_{\min})} + \frac{1}{2} \frac{(1+2\beta y_{\min}/\alpha)(\sqrt{\alpha} \Delta t E - y_{\min})}{\sqrt{s(y_{\min})}} \\ \text{near } \sqrt{\alpha} \Delta t E = y_{\min} \\ \sqrt{s(y_{\max})} + \frac{1}{2} \frac{(1+2\beta y_{\max}/\alpha)(\sqrt{\alpha} \Delta t E - y_{\max})}{\sqrt{s(y_{\max})}} \\ \text{near } \sqrt{\alpha} \Delta t E = y_{\max} \end{cases}$$

The sought region of the dynamic range can then be written as:

$$\begin{aligned} y'_{\min} < \sqrt{\alpha} \Delta t E < y'_{\max} \quad \text{with} \\ y'_{\min} &= \frac{k(\alpha s(y_{\min}) - \beta y_{\min}^2) + y_{\min} \alpha (\sqrt{s(y_{\min})} - k/2)}{\alpha \sqrt{s(y_{\min})} - k(\sqrt{\alpha} + 2\beta y_{\min})/2} \quad \text{and} \\ y'_{\max} &= \frac{-k(\alpha s(y_{\max}) - \beta y_{\max}^2) + y_{\max} \alpha (\sqrt{s(y_{\max})} + k/2)}{\alpha \sqrt{s(y_{\max})} + k(\sqrt{\alpha} + 2\beta y_{\max})/2} \end{aligned} \quad (48)$$

An illustration of the bounds as a function of the offset noise level $\sigma_{\alpha_j}^2$ and the amplification factor $\sqrt{\alpha}$ is given in Figure 15. It can be seen that even at low SNRs, the region $[y'_{\min}, y'_{\max}]$ still covers a sufficiently large region near the center of the dynamic range $\sqrt{\alpha} \Delta t E = (y_{\min} + y_{\max})/2$.

Appendix 3: optimization problem for determining the weighting function

In this section, we will solve the following optimization problem:

$$\min_{\vec{w}} \frac{\sum_{j=1}^P w_j^2 \alpha_j^2 \sigma_j^2}{\left(\sum_{j=1}^P w_j \alpha_j\right)^2} \quad \text{such that } \|\vec{w}\|_{\infty} = 1, \quad (49)$$

where $\|\cdot\|_{\infty}$ is the maximum norm. To deal with the constraint $\|\vec{w}\|_{\infty} = 1$, we introduce a new constraint $\sum_{j=1}^P w_j \alpha_j = C$ and subsequently we fix the constant C by imposing that $\|\vec{w}\|_{\infty} = 1$. Using Lagrange multipliers, the resulting problem can be stated as follows:

$$\min_{\vec{w}} \sum_{j=1}^P w_j^2 \alpha_j^2 \sigma_j^2 + \lambda \left(\sum_{j=1}^P w_j \alpha_j - C \right). \quad (50)$$

Taking the derivatives of the cost function with respect to w_j gives the following equations:

$$w_j \alpha_j \sigma_j^2 = -\lambda/2, \quad \text{with } j = 1, \dots, P.$$

One of the issues here is that σ_j^2 may become zero, such that the weight w_j can not be determined. In this case, we are free to choose the weight, and for convenience

we choose $w_j = 1$ (such that the contribution of the corresponding LDR in the HDR reconstruction will be *maximal*). Imposing the constraints $\sum_{j=1}^P w_j \alpha_j = C$ and $\|\vec{w}\|_{\infty} = 1$ then leads to the following solution:

$$w_j = \begin{cases} 1 & \sigma_j^2 = 0 \\ \frac{\sigma_j^{-2}}{\alpha_j} / \max_{\sigma_j^2 \neq 0} \frac{\sigma_j^{-2}}{\alpha_j} & \sigma_j^2 \neq 0 \end{cases}$$

This solution is used in Section “Probabilistic formulation of the weighting function” for computing optimal weighting functions.

Appendix 4: derivation of a direct formula for the weighting function

Based on the noise model from Equations (17) and (30), the PDF $f_{\gamma^{-1}(z)|E}(\gamma^{-1}(z) | E_i)$ can be determined as follows:

$$\begin{aligned} f_{\gamma^{-1}(z)|E}(\gamma^{-1}(z) | E_i) &= \frac{1}{\sqrt{2\pi s'(E_i)}} \exp \\ &\left(-\frac{1}{2} \left(\frac{\gamma^{-1}(z) - \sqrt{\alpha_j} \Delta t_j E_i}{s'} \right)^2 \right) \left| \frac{d\gamma^{-1}(z)}{dz} \right| \end{aligned} \quad (51)$$

with $s' = \text{Var}[\gamma^{-1}(z) | E_i]$. Using results from Appendix 1, we approximate $s'(E_i)$ as follows:

$$s' \approx s(\gamma^{-1}(z)) = \sigma_{\alpha_j}^2 + \alpha_j \gamma^{-1}(z) + \beta_{ij} (\gamma^{-1}(z))^2. \quad (52)$$

The weighting function (32) can be computed using Bayes' rule:

$$w_{ij}^{(2)} = \frac{\int_{E_{\min}^{(ij)}}^{E_{\max}^{(ij)}} f_{\gamma^{-1}(z)|E}(\gamma^{-1}(z) | E_i) f_E(E_i) dE_i}{\int_{-\infty}^{+\infty} f_{\gamma^{-1}(z)|E}(\gamma^{-1}(z) | E_i) f_E(E_i) dE_i} \quad (53)$$

Using (51) and the approximation (52), this weighting function amounts to the direct expression:

$$\begin{aligned} w_{ij}^{(2)} &\approx \frac{\Phi(y'_{\min}(i,j)) - \Phi(y'_{\max}(i,j))}{\Phi(0) - 1} \quad \text{with } : \Phi(x) \\ &= \text{erf} \left(\frac{x - \gamma^{-1}(z_{ij}) + \lambda \alpha_j^{-1/2} \Delta t_j^{-1} s(\gamma^{-1}(z_{ij}))}{\sqrt{s(\gamma^{-1}(z_{ij}))}} \right). \end{aligned}$$

Endnotes

^a We remark that due to quantum-mechanical aspects, in reality there is no such thing as an “ideal” signal, here we define the “ideal” signal as the voltage averaged over a “very” long exposure time.

^b Remark that Bayesian estimates (such as MAP, MMSE) are equally possible, given that prior information on E_i is available.

^c Due to Jensen's inequality.

^d This is because the contribution of several noise sources can be modeled using a Gaussian distribution, with signal-dependent mean and variance.

^e Equation (13) is invariant under scaling of w_{ij} ; if a certain w_{ij}^* minimizes MSE then any scaled version aw_{ij}^* with $a > 0$ is also a solution to (24). By adding an extra constraint to the solution, this problem is solved.

^f More specifically, we use $\left| \sum_j a_j b_j \right|^2 < \sum_j a_j^2 \sum_l b_l^2$ with $a_i = w_{ij} \sigma^{-2} (E_i, \Delta t_j) v (E_i, \Delta t_j)$ and $b_l = 1$.

^g This follows from the fact that $E_i > 0$, $\Delta t_j > 0$, $\alpha_j > 0$ and $\sigma_{0,j}^2 > 0$.

^h The exposure time is either stored in the RAW data files, or the EXIF information of the compressed JPEG files.

ⁱ To see this, note that $\sqrt{s(\Delta t \alpha)} / (\sqrt{\Delta t \alpha}) = \sqrt{\frac{\sigma_{\alpha}^2}{\alpha \Delta t^2} + \Delta t E + \frac{\beta}{\alpha} E}$, which increases monotonically in σ_{α}^2 and β and decreases in α .

^j Available at <http://people.csail.mit.edu/hasinoff/hdrnoise/>.

^k <http://rawtherapee.com/>.

Competing interests

The authors declare that they have no competing interests.

Acknowledgements

The authors thank Dr. Filip Rooms and Koen Douterloigne for providing HDR datasets. Bart Goossens acknowledges the financial support from the Flemish Research Foundation (FWO), Belgium.

Received: 29 June 2011 Accepted: 14 May 2012

Published: 16 August 2012

References

- DL Donoho, De-noising by soft-thresholding. *IEEE Trans. Inf. Theory*. **41**(3), 613–627 (1995)
- RR Coifman, DL Donoho, in *Wavelets and Statistics*, ed. by A Antoniadis, G Oppenheim. Translation-Invariant Denoising. (Springer Verlag, New York, 1995), pp 125–150
- S Chang, B Yu, M Vetterli, Spatially adaptive wavelet thresholding with context modeling for image denoising. *IEEE Trans. Image Process.* **9**(9), 1522–1531 (2000)
- J Portilla, V Strela, M Wainwright, E Simoncelli, Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Trans. Image Process.* **12**(11), 1338–1351 (2003)
- A Pižurica, W Philips, Estimating the probability of the presence of a signal of interest in multiresolution single- and multiband image denoising. *IEEE Trans. Image Process.* **15**(3), 654–665 (2006)
- LI Rudin, S Osher, E Fatemi, Nonlinear total variation based noise removal algorithms. *Physica D*. **60**, 259–268 (1992)
- P Perona, J Malik, Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(7), 629–639 (1990)
- J Weickert, *Anisotropic Diffusion in Image Processing*. ser. ECVI Series (Teubner-Verlag, Stuttgart, 1998)
- A Buades, B Coll, J Morel, A review of image denoising algorithms, with a new one. *SIAM Interdiscip. J. Multiscale Model. Simu.* **4**(2), 490–530 (2005)
- C Kervrann, J Boulanger, Optimal spatial adaptation for patch-based image denoising. *IEEE Trans. Image Process.* **15**(10), 2866–2878 (2006)
- B Goossens, J Aelterman, HQ Luong, A Pižurica, W Philips, in *2009 International Workshop on Local and Non-Local Approximation in Image Processing (LNLA2009)*. Efficient design of a low redundant discrete shearlet transform (Tuusula, Finland, Aug 2009), pp. 112–124
- S Lim, in *Proc. SPIE Digital Photography II*, vol. 6069. Characterization of noise in digital photographs for image processing (San José, CA, USA, 2006), p. p. 606900
- K Hirakawa, in *Color Filter Array Image Analysis for Joint Denoising and Demosaicking* ed. by R Lukac. Single-Sensor Imaging: Methods and Applications for Digital Cameras. (CRC Press, Boca Raton, 2008)
- FJ Anscombe, The transformation of Poisson, binomial and negative-binomial data. *Biometrika*. **35**, 245–254 (1948)
- M Makitalo, A Foi, Optimal inversion of the anscombe transformation in low-count poisson image denoising. *IEEE Trans. Image Process.* **20**(1), 99–109 (2011)
- H Faraji, WJ MacLean, CCD noise removal in digital images. *IEEE Trans. Image Process.* **15**(9), 2676–2685 (2006)
- C Liu, R Szeliski, SB Kang, CL Zitnick, WT Freeman, Automatic estimation and removal of noise from a single image. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(2), 299–314 (2008)
- A Foi, M Trimeche, V Katkovnik, K Egiazarian, Practical Poissonian–Gaussian noise modeling and fitting for single-image raw-data. *IEEE Trans. Image Process.* **17**(10), 1737–1754 (2008)
- A Foi, S Alenius, M Trimeche, V Katkovnik, K Egiazarian, in *Proc. IEEE International Conference on Image Processing ICIP 2005*, A spatially adaptive Poissonian image deblurring, vol. 1, 11–14 Sept 2005, pp. I-925–8
- D Paliy, A Foi, R Bilcu, V Katkovnik, K Egiazarian, in *Proceedings of European Signal Processing Conference Joint deblurring and demosaicing of Poissonian bayer-data based on local adaptivity*, (Lausanne, 2008), pp. 1569104966/1–5
- DT Kuan, A Sawchuk, TC Strand, P Chavel, Adaptive noise smoothing filter for images with signal-dependent noise. *IEEE Trans. Pattern Anal. Mach. Intell.* **7**(2), 165–177 (1985)
- F Argenti, G Torricelli, L Alparone, in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2002)*, Signal-dependent noise removal in the undecimated wavelet domain, vol. 4, 13–17 May 2002, pp. IV-3293–IV-3296
- B Goossens, A Pižurica, W Philips, in *IEEE International Conference on Image Processing (ICIP) Wavelet domain image denoising for non-stationary and signal-dependent noise*, (Atlanta, 2006), pp. 1425–1428
- SW Hasinoff, F Durand, WT Freeman, in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR) Noise-optimal capture for high dynamic range photography*, (San Francisco, 2010), pp. 553–560
- K Hirakawa, PJ Wolfe, in *Proc. 17th IEEE Int Image Processing (ICIP) Conf* Optimal exposure control for high dynamic range imaging, (Hong Kong, 2010), pp. 3137–3140
- G Larson, H Rushmeier, C Piatko, A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. Vis. Comput. Graph.* **3**, 291–306 (1997)
- S Mann, R Picard, in *In Proceedings of IS&T 48th annual conference (Cambridge, Massachusetts, May 1995)* On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures, (Los Angeles, 1994), pp. 422–428
- P Debevec, J Malik, in *Proceedings of SIGGRAPH97, Computer Graphics Proceedings Recovering high dynamic range radiance maps from photographs*, (Ft. Collins, 1997), pp. 369–378
- T Mitsunaga, S Nayar, Radiometric self calibration. *IEEE Conf. Comput. Vision and Pattern Recognit (CVPR)*. **1**, 374–380 (1999)
- S Mann, Comparametric equations with practical applications in quantigraphic image processing. *IEEE Trans. Image Process.* **9**(8), 1389–1406 (2000)
- S De Neve, B Goossens, H Luong, W Philips, in *IEEE Int. Conf. Image Processing (ICIP2009)* An improved HDR image synthesis algorithm, (Cairo, 2009), pp. 1545–1548
- E Reinhard, G Ward, S Pattanaik, P Debevec, *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting* (Morgan Kaufmann Publishers, San Francisco, 2005)
- Y Tsin, V Ramesh, T Kanade, in *Proc. IEEE Int. Conf. Comp. Vis.* Statistical calibration of the CCD imaging process, (Vancouver, 2001), pp. 480–487
- K Kirk, HJ Andersen, in *Proc. British Mach. Vis. Conf.* Noise characterization of weighting schemes for combination of multiple exposures, (Edinburgh, 2006), pp. 1129–1138
- A Foi, Clipped noisy images: heteroskedastic modeling and practical denoising. *Signal Process.* **89**(12), 2609–2629 (2009)
- S Lim, A El Gamal, Gain fixed pattern noise correction via optical flow. *IEEE Trans. Circ. Syst. I Fundament. Theory Appl.* **51**(4), 779–786 (2004)
- MD Grossberg, SK Nayar, Modeling the space of camera response functions. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(10), 1272–1282 (2004)

38. M Robertson, S Borman, R Stevenson, Estimation-theoretic approach to dynamic range improvement using multiple exposure. *J. Electr. Image.* **12**(2), 219–228 (2003)
39. S Lin, J Gu, S Yamazaki, Shum H-Y, in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition CVPR 2004*, vol. 2, Radiometric calibration from a single image, (Washington, 2004), pp. II-938 - II-945
40. GH Golub, CF Van Loan, *Matrix Computations*, 3rd edn. (Johns, Hopkins University Press, Baltimore, MD, USA, 1996)
41. P Gill, W Murray, MH Wright, *Practical Optimization* (Academic Press, London, UK, 1981)
42. J Portilla, in *Proc. of the SPIE's 50th Annual Meeting: Wavelets XI*, Image restoration using Gaussian scale mixtures in overcomplete oriented pyramids (a review), vol. 5914, (San Diego, CA, Aug 20050, pp, 468–482

doi:10.1186/1687-6180-2012-171

Cite this article as: Goossens et al.: Realistic camera noise modeling with application to improved HDR synthesis. *EURASIP Journal on Advances in Signal Processing* 2012 **2012**:171.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Immediate publication on acceptance
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ springeropen.com
