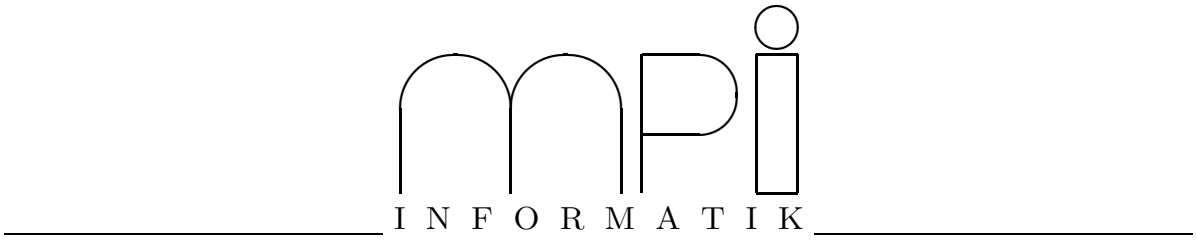

Photometric Calibration of High Dynamic Range Cameras

Grzegorz Krawczyk Michael Goesele
Hans-Peter Seidel

MPI–I–2005–4–005 May 2005

FORSCHUNGSBERICHT RESEARCH REPORT

MAX-PLANCK-INSTITUT
FÜR
INFORMATIK

Stuhlsatzenhausweg 85 66123 Saarbrücken Germany

Authors' Addresses

Grzegorz Krawczyk
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
krawczyk@mpi-inf.mpg.de

Michael Goesele
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
goesele@mpi-inf.mpg.de

Hans-Peter Seidel
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
hpseidel@mpi-inf.mpg.de

Abstract

The current trends in computer graphics focus on acquiring, processing and reproducing real world scenes, often with a quite high dynamic range (HDR). Many algorithms require such HDR image or video data as an input and the increasing availability of HDR sensors provides a natural source of them. But at the same time the diversity of HDR sensor technology leads to a strong need for photometric calibration of such HDR camera systems in order to retain the fidelity of the final result.

We present a calibration approach based on an earlier HDR recovery algorithm for standard low dynamic range (LDR) cameras. We focus here on the specific challenges posed by cameras with a dynamic range of more than six orders of magnitude such as complex camera response curves or selection of appropriate calibration targets. We perform an absolute calibration of two HDR and one LDR camera systems to allow for recovery of real-world luminance values. To validate our approach, we compare these luminance values to measurements performed with a luminance meter for grayscale patches covering a dynamic range of six orders of magnitude. The achieved accuracy of the photometric calibration is sufficient for many measurement and image-based acquisition applications.

Keywords

calibration, photometry, high dynamic range, video camera, absolute calibration

1 Introduction

High dynamic range (HDR) imaging technology is becoming a standard tool in computer graphics and vision applications. Especially when acquiring images in complex lighting situations, standard low dynamic range (LDR) imaging sensors (mostly based on CCD technology [4]) are easily saturated and image details are lost. Many applications such as measuring material characteristics, capturing environment maps for realistic rendering using global illumination solutions, or image-based recovery of surface reflection parameters furthermore require photometrically calibrated images with absolute luminance values per pixel.

Several methods to improve the dynamic range of LDR cameras have been developed based on the multiple acquisition of the target scene with different exposure settings [1, 8, 9, 2]. These methods, however, impose a severe limitation on the applications as a precise HDR acquisition is only possible for static scenes. Kang et al. [5] proposed a method for dynamic scenes but the achievable dynamic range is limited.

In recent years, several new image sensors, mostly based on CMOS technology, have been developed that are capable of capturing images with a dynamic range of up to 8 orders of magnitude at video frame rates. Several HDR video cameras based on these sensors are already commercially available (see [5] for an overview). They cannot only be used to speed up many image-based measurement systems that currently use multi-exposure techniques but allow also to capture dynamic scenes with high contrast. Compared to multi-exposure approaches, HDR video cameras offer considerably wider dynamic range and the quality is independent of changes in scene content as frame-to-frame coherence is not required.

Although photometric calibration of LDR cameras has already been widely discussed, the equivalent task for HDR cameras is not trivial. The technologies developed to capture high dynamic range are varied, so that it is difficult to define a general model of camera response curves. Depending on their built-in assumptions many algorithms suited for LDR cameras may fail. Furthermore, finding a suitable calibration target with sufficiently high dynamic range is difficult. Finally, it is often not possible to control the integration time of HDR cameras so that acquiring differently exposed images requires additional optical filters.

1.1 Contribution

The goal of this paper is to perform absolute photometric calibration of HDR cameras and to validate the accuracy of current HDR cameras for

applications requiring such calibration. We first adapt an existing technique by Robertson et al. [9] to the specific requirements of HDR camera systems. Alternatively, we follow earlier approaches [7] and make an assumption about the shape of the response curve based on the used sensor technology and camera data sheets. We thus only have to determine basic offset and scale values to perform an absolute photometric calibration. We then compare absolute luminance values of a series of grayscale patches determined using the calibrated cameras with ground truth measurements performed with a luminance meter and discuss the error in the light of possible applications.

The remainder of the paper is structured as follows: after an overview over previous work in Section 2, we discuss our methods for photometric calibration of HDR cameras (Section 3). In Section 4, we present experimental results to validate our approach. Before concluding, we discuss the benefits and drawbacks of the approach in Section 5.

2 Previous Work

We first give a brief overview over current HDR camera technology. We then present existing methods for photometric calibration of LDR cameras and discuss their applicability to HDR camera systems.

2.1 HDR Camera Technology

There are currently two major approaches to extend the dynamic range of an imaging sensor. One type of sensor collects charge generated by the photo current. The amount of charge collected per unit of time is linearly related to the irradiance on the chip (similar to a standard CCD chip [4]), the exposure time is however varying per pixel (sometimes called "locally autoadaptive" [6]). This can for instance be achieved by sequentially capturing multiple exposures with different exposure times setting or by stopping after some time the exposure of the pixels that would be overexposed during the next time step. A second type of sensor uses the logarithmic response of a component to compute the logarithm of the irradiance in the analog domain. Both types require a suitable analog-digital conversion and generate typically a non-linearly sampled signal encoded using 8–16 bits per pixel value.

2.2 Photometric Calibration

Photometric calibration of LDR cameras has been widely discussed in the context of high dynamic range imaging. A good overview over existing work

can be found in [9]. Current methods to extend the dynamic range through multiple exposures consist of two steps: they first recover the camera response curve in order to linearize the captured images. If the sensor has a linear characteristic (such as the raw output of a CCD sensor) this step can be omitted [7]. In a second step, the linearized input images are scaled according to the exposure settings and combined into a single image.

In principle, three different approaches are used for response curve recovery of LDR cameras. We briefly discuss each of them in view of possible application to photometric calibration of HDR cameras.

The algorithm developed by Debevec and Malik [1] is based on the concept that a particular pixel exposure is defined as a product of the irradiance at the film and the exposure time, transferred by the camera response function. This concept is embedded in an objective function which is minimized to determine the camera response curve. The objective function is additionally constrained by the assumption that the response curve is smooth, which is essential for the minimization process. Whereas this assumption is generally true for LDR cameras based on CCD technology, the response curve is normally not smooth in locally autoadaptive HDR sensors. Furthermore, the process of recovering the response curve is based on solving a set of linear equations. While the size of the matrix representing these linear equations is reasonable for 8 bit data, memory problems may occur for arbitrary precision data typical to HDR acquisition so that extensive subsampling is required.

The method proposed by Mitsunaga and Nayar [8] computes a radiometric response function approximated using a high-order polynomial without precise knowledge of the exposures used. The recovery process is limited to estimating the order of the polynomial and its coefficients. The authors state that it is possible to represent virtually any response curve using a polynomial. This fact is true for LDR cameras based on a CCD sensor, however it is not possible to approximate the logarithmic response of some CMOS sensors in this manner. Polynomial approximation also assumes that the response curve is continuous, which depends on the encoding.

The response curve recovery proposed by Robertson et al. [9] is based on an estimation approach. The authors define an observation model which assumes a linear relation between the camera response values and a product of irradiance and exposure time. The response curve is determined through an iterative algorithm minimizing the objective function independently for each camera output level, making no assumption on the shape of the response curve. The memory requirements are limited to the storage of the source exposures and no additional space is necessary for the computations.

3 HDR Camera Response

To choose an appropriate algorithm for the recovery of the HDR camera response, we can formulate the following requirements. First, varied technology used in the HDR sensors results in non-linear, often non-continuous output. Therefore no assumption on the response curve shape can be made. Second, the range of HDR camera output levels is in general arbitrary and often exceeds 8 bits. Thus the algorithm should be well scalable in terms of memory consumption and computational complexity for an arbitrary number of camera output levels. In view of the analysis given in the previous section, the calibration approach by Robertson et al. [9] appears to be the most suitable method for HDR cameras. In the following section we briefly summarize their algorithm.

3.1 Response Recovery

The camera response curve is derived from an observation model describing the relation between the camera output levels produced under known exposure to the underlying light intensity values. If the camera output value for j^{th} pixel of the i^{th} exposed image is denoted by y_{ij}, the exposure of this image is denoted by t_i, and the underlying light values by x_j, the observation model can be described as follows. If only the exposure is varying between images, the amount of light contributing to the output value y_{ij} is $t_i * x_j$, with an additional noise term N_{ij}^c. Therefore the amount of the captured light plus noise is transferred by the camera response function f to the exposed image in the following way:

$$y_{ij} = f(t_i * x_j + N_{ij}^c) \tag{1}$$

Based on the above observation, the algorithm computes the inverse of the response curve f^{-1} by minimizing the following objective function $\tilde{O}$:

$$\tilde{O}(f, x) = \sum_{i,j} w_{ij} * (f^{-1}(y_{ij}) - t_i * x_j) \tag{2}$$

The noise term in Equation 1 is now represented by the weighting function w_{ij} which models the confidence in the accuracy of the values captured by the camera. The function w_{ij} is a Gaussian distributed around the medium output value. It implies that values located near the minimum or maximum are acquired with less confidence than pixel values in the middle range of the camera response.

The objective function $\tilde{O}$ is numerically minimized using a form of Gauss-Seidel relaxation. For a color camera, the response curve should be determined for each color channel separately unless it is known that the color

channels are treated exactly identical. For a detailed description of this algorithm we refer to [9].

3.2 Target Setup

For the accuracy and completeness of the response curve, it is necessary that the pixel values of the acquired exposures uniformly cover the full range of the camera output values. Due to the fact that the camera captures high dynamic range, a proper scene setup becomes an issue. In fact we were not able to prepare in our lab a single setup that satisfies this requirement. It is therefore necessary to capture several different setups, each covering a different luminance range, and to stitch them together into a single image. With this approach it is in principle possible to provide data of arbitrarily high dynamic range as input to the response recovery algorithm.

The estimation of the camera response curve is based on different exposures of the same scene. Using conventional cameras, different exposures are acquired by changing either aperture size or integration time. In the case of an HDR camera, changing the integration time is however not possible. It is furthermore desirable to keep the aperture constant to prevent artifacts related to varying depth of field, vignetting, and to diffraction effects at the aperture blades. We therefore simulate different exposures by mounting neutral density filters in front of the camera lens. Following the analysis in [2], we suggest to acquire two images that are exposed similarly and one that is considerably different. If possible, the properties of the sensor technology should also be taken into account: if a camera varies exposure times in powers of 2, the filters' extension factors should not be powers of 2 to avoid mapping possible singularities onto each other.

The calibration of video cameras allows also easily to capture a larger number of frames for all input scenes in order to reduce the influence of image noise on the response curve recovery.

3.3 Absolute Calibration

The relative luminance values obtained from the recovered response curve are linearly proportional to the absolute luminance with a scale factor dependent on the exposure parameters and the lens system. Absolute calibration is based on the acquisition of a scene containing patches with known reflectance. The scale factor is determined by minimization of the relative error between measured luminance values and the captured, linearized camera response.

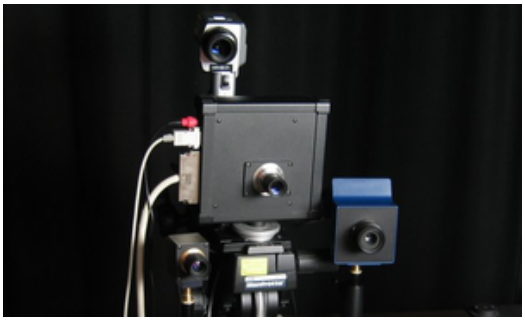

Figure 1: Cameras used in our experiment: HDRC VGAx camera (lower left), Silicon Vision Lars III (center), Jenoptik C14 (lower right), and Minolta LS-100 luminance meter (top).

4 Example Calibration

In this section, we validate the proposed process of photometric calibration by performing an example calibration of two HDR video cameras: a Silicon Vision Lars III and a HDRC VGAx camera. For comparison purposes we also included in the experiment the Jenoptik C14 – a high-end, CCD based LDR camera (see Figure 1).

The Lars III sensor is an example of a locally autoadaptive image sensor [6]: the exposure is terminated for each individual pixel after one out of 12 possible exposure times (usually powers of 2). For every pixel, the camera returns the amount of charge collected until the exposure was terminated as a 12 bit value and a timestamp. This can be lossless encoded as 16 bit value similar to an unsigned IEEE floating point value [3] with 4 bit exponent e and 12 bit mantissa m. A linearized relative luminance value can also be reconstructed as $m \cdot 2^e$.

The HDRC sensor is a logarithmic-type sensor [10]. The camera returns images of 10 bit values.

4.1 Experiment Setup

In order to cover the dynamic range in which these cameras operate, we acquired three target setups with varied luminance characteristic (see Figure 2). The target scene was first acquired under a moderate illumination covering lower parts of the dynamic range. Then a strong light source was included into the setup providing both strong illumination and hard shadows. To provide sufficient samples in the upper dynamic range, we acquired the light source with reflector shining directly towards the cameras. Stitching these

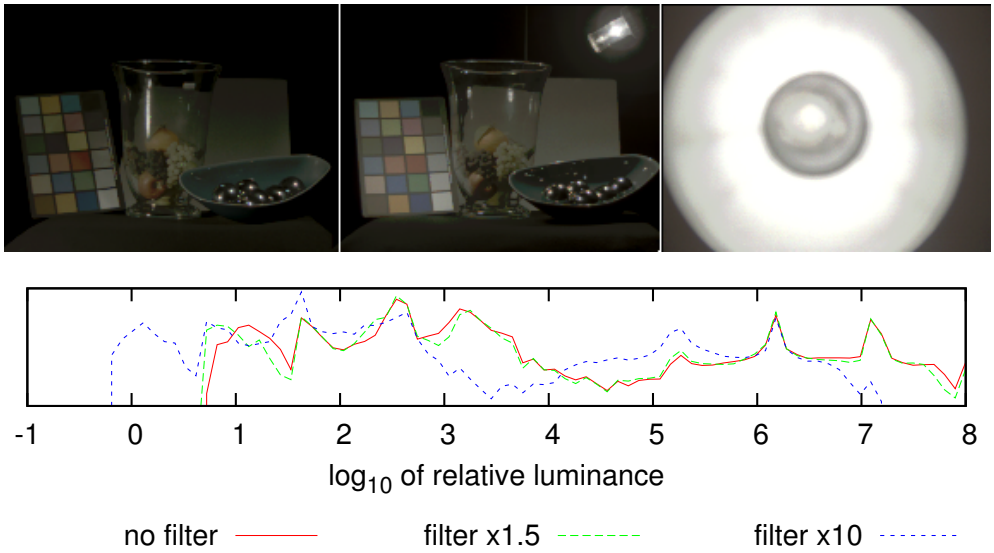

Figure 2: Three target setups used for the measurement of the response curve (tone mapped for presentation). The histogram shows the luminance distribution in the stitched images for acquisition without filter, and using a neutral density filter with $\times 1.5$ and $\times 10$ extension factors. This setup covers 8 orders of luminance magnitude.

three images together yields an input for the response recovery algorithm covering a dynamic range of more than 8 orders of magnitude. To obtain differently exposed images, each target setup was acquired without any filter and with neutral density filters with $\times 1.5$ and $\times 10$ extension factor. The response curve for the LDR camera was recovered using a series of 13 differently exposed images of a GretagMacbeth ColorChecker and a standard algorithm [9].

4.2 Recovered Response Curves

The recovered response curves for the three cameras are shown in Figure 3. A single response curve was recovered for the monochromatic Lars III camera whereas separate curves were determined for the three color channels of the other cameras. As we had access to the raw sensor values of the HDRC camera before Bayer interpolation, we recovered the response curve for each channel directly from corresponding pixels in order to avoid possible interpolation artifacts.

Figure 3 shows that the response curves of the two HDR cameras both

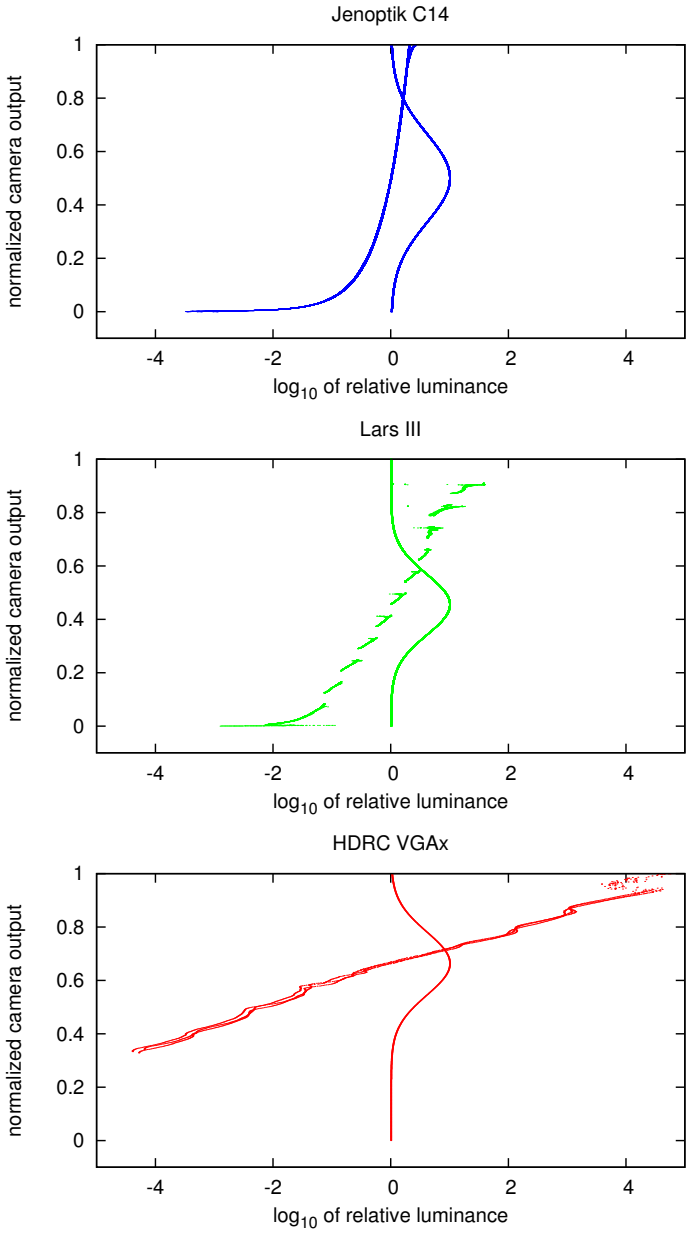

Figure 3: The recovered response curves and corresponding weighting functions representing the confidence in the accuracy of the acquisition for given camera output value (value 1.0 represents the full confidence, 0.0 represents no confidence).

cover a considerably wider range of luminance than the high-end LDR camera that covers a range of about 3.5 orders of magnitude. The different shapes of the HDR response curves are caused by their respective sensor technology and the encoding.

The logarithmic HDRC VGAx camera has the highest dynamic range (more than 8 orders of magnitude). Unfortunately, an offset in the A/D conversion makes the lower third of the 10 bit range unusable.

The multiple exposure values of the locally autoadaptive Lars III camera are well visible as discontinuities in the response curve. Note that the luminance range is covered continuously and gaps are only caused by the encoding. The camera covers a dynamic range of about 5 orders of magnitude. Noise at the switching points between exposure times is well visible.

4.3 Results of Absolute Calibration

The recovered response curves convert the camera output values into relative luminance values. To perform an absolute calibration, we acquired a GretagMacbeth ColorChecker chart under 6 different illumination conditions. The luminance of the gray patches was measured using a Minolta LS-100 luminance meter yielding a total of 36 samples and an optimal scale factor was determined for each camera. Results for selected patches are shown in Table 1.

The accuracy of the absolute calibration for the 36 patches can be seen in Figure 4. The calibrated camera luminance values are well aligned to the measured values proving that the response curve recovery was accurate. To quantify the quality of the absolute calibration, we calculated the average relative error for these data points. For the HDRC camera, relative error in the luminance range of 1–10,000 $[cd/m^2]$ is 13% while the relative error for the Lars III camera in the luminance range of 10–1,000 $[cd/m^2]$ amounts to 9.5%. Note that these results can be obtained with a single acquisition. Using multiple exposures, the C14 camera is capable of an average relative error of below 7% in the range 0.1–25,000 $[cd/m^2]$, thus giving the most accurate results.

4.4 Fitting a Response Function

As an alternative approach to the previous sections, we fit the camera output values directly to the measured luminance values from the experiment in Section 4.3 using an a priori function appropriate for the given HDR sensor. Thus, for the HDRC camera we fit the parameters of a logarithmic function $y_j = a * \log_{10}(x_j) + b$ and for the decoded values of the Lars III camera we

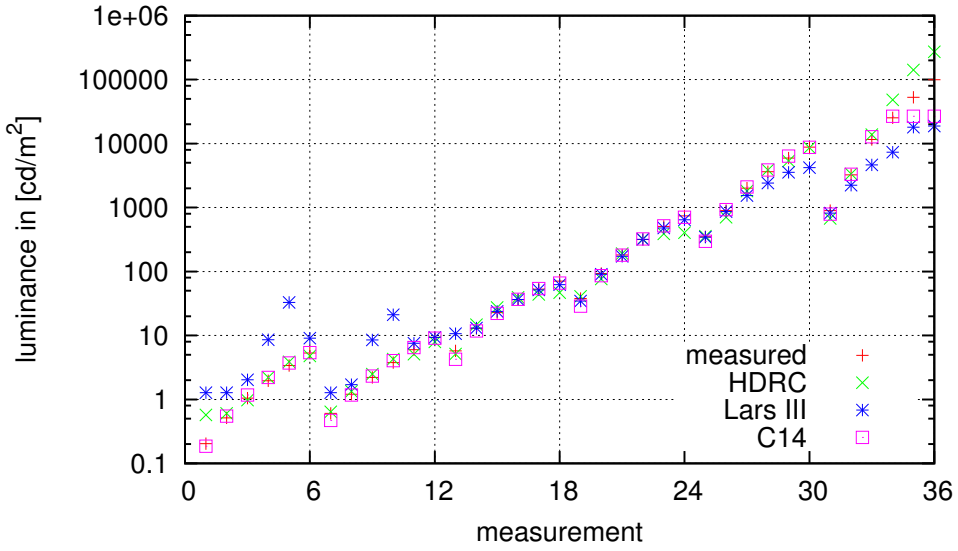

Figure 4: The results of absolute calibration. The recovered response curves were fitted to the measurements of 6 gray patches of GretagMacbeth ColorChecker chart under 6 different illumination conditions.

fit a linear function $y_j = a * x_j + b$. Results for selected patches are shown in Table 1.

We compare the relative errors achieved by the response curve recovery including absolute calibration and the function fit in Figure 5. The average relative error is about 6% for the HDRC camera and luminance values above $1\,[cd/m^2]$. For the Lars III camera it is also about 6% for luminance values above $10\,[cd/m^2]$. Especially for high luminance values above $10,000\,[cd/m^2]$, the calibration via function fitting provides more accurate results.

In addition, the fitting approach allows to extrapolate the response curves for values beyond the range of the calibration targets. To verify this, we acquired an extremely bright patch ($194,600\,[cd/m^2]$) and compared the measurement of the light meter to the calibrated response of the HDR cameras. As shown in Table 1, only the readout from the HDRC camera derived via function fitting is reliable while the HDRC response curve seems to be bogus in that luminance range. The Lars III camera reached the saturation level and yielded arbitrary results (see Table 1). Likewise, this patch could not be recorded with the available settings of the LDR camera.

measured luminance $[cd/m^2]$	HDRC output 10 bit	HDRC recovered resp. $[cd/m^2]$	HDRC fitted resp. $[cd/m^2]$	Lars III output 16 bit	Lars III recovered resp. $[cd/m^2]$	Lars III fitted resp. $[cd/m^2]$	C14 recovered resp. $[cd/m^2]$
5.3	404	4.78	4.57	200	9.04	8.69	5.38
9.3	424	7.95	8.12	366	9.35	11.52	9.05
70.9	497	46.70	62.65	3,278	62.46	61.33	66.07
741.2	582	403.22	695.22	18,722	643.96	663.00	704.43
8,796	673	8,616.82	8,924.89	32,126	4,196.72	7,822.66	8,734.86
194,600	788	1,081,800	225,010	43,665	44,121,000	50,415	-

Table 1: Measured luminance values of patches with corresponding camera output and calibrated camera response obtained via response recovery (Sections 4.2 and 4.3) and function fitting (Section 4.4).

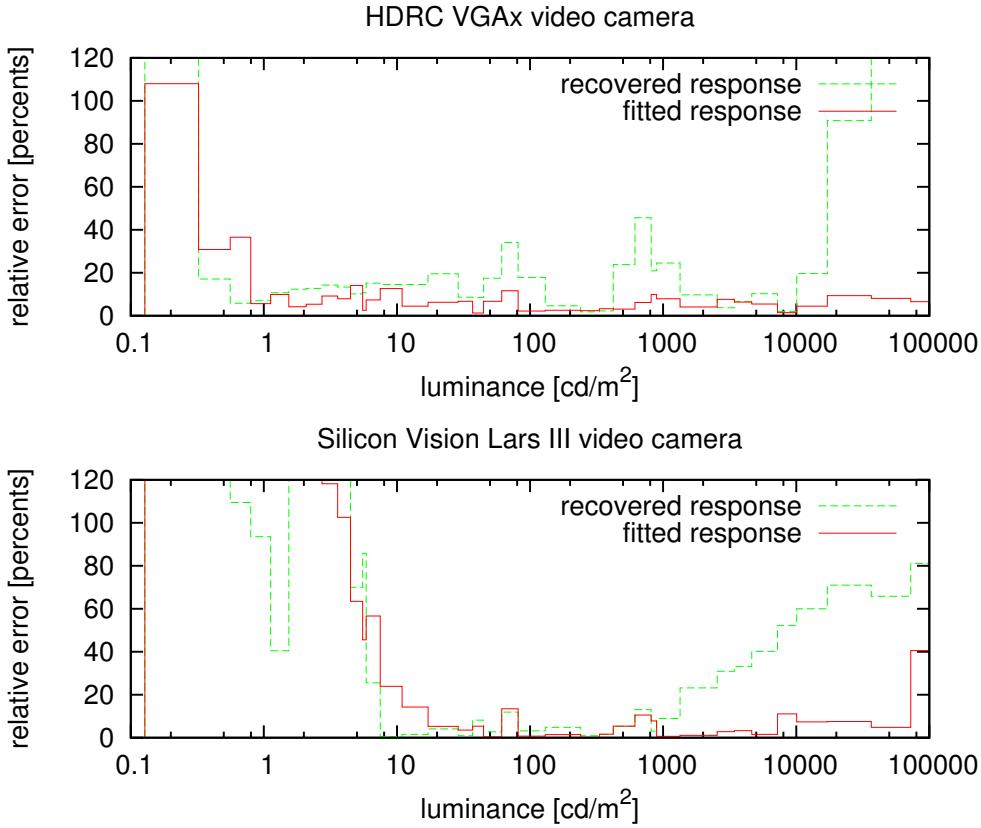

Figure 5: Comparison of the relative errors in luminance measurement achieved by the response curve recovery including absolute calibration and by the function fit.

5 Discussion

The ability to capture HDR data has a strong impact on various computer graphics applications. The flexibility in the acquisition of dynamic sequences, which can contain both bright light sources and very dark areas at the same moment (such as sun and deep shadows), is unprecedented. Dynamic environment maps can for instance be captured in real time to faithfully convey the illumination conditions of the real world to rendering algorithms, thus giving new insights to mixed reality applications. Predictive rendering methods may additionally benefit from knowledge about physically accurate luminance values, so that e.g. the results of global illumination solutions can be directly compared to real world measurements. Also, perceptually enabled algorithms may appropriately simulate the behavior of the human visual

system, since precise luminance values can be captured. The application of HDR cameras is natural in such situations but requires careful calibration: Computer graphics algorithms commonly operate in linear intensity space. Understanding and calibrating the camera response is crucial for the proper interpretation of captured data.

In this paper, we presented two alternative approaches to photometric calibration of HDR cameras. Although the relative error achieved by the function fitting approach is lower, the response recovery algorithm is useful to obtain the exact shape of the camera response curve and to give confidence that the chosen a priori function is correct. It can also help to understand the behavior of the sensor, especially if the encoding is unknown. The low precision of the measurements in the luminance range above $1,000\,[cd/m^2]$ is a clear limitation which could be caused either by an excessive amount of noise or by glare artifacts introduced by the ND filters.

The function fitting approach has strong advantages in the quality of the results and the ability to extrapolate from the calibration data. The confidence in extrapolated measurements is however limited and the error cannot be predicted because the exact shape of the response function in this range is unknown. Nevertheless, it is important to keep in mind that the fitting approach led to a lower relative error for both HDR cameras.

The experiments showed also that the HDR cameras have problems in dark conditions below $10\,[cd/m^2]$ which can be explained by the high noise level in the sensors. The quality of a high-end CCD camera such as the Jenoptik C14 combined with traditional HDR recovery algorithms still cannot be achieved consistently over the whole dynamic range of the HDR cameras.

Finally, the accuracy of the photometric calibration is not the only important quality measure. Depending on the application, other issues such as the quantization of the luminance values might have an important influence on the quality of the measurements and need to be further investigated.

6 Conclusions

Driven by vision applications in difficult lighting situations, the evolution of high dynamic range camera systems will continue. We showed in this paper that such cameras can also be used for photometric measurements with an accuracy comparable to traditional multi-exposure HDR recovery systems although some issues such as the noise level in dark conditions still need to be resolved. We expect that many future applications in computer graphics will especially benefit from the fast acquisition of high dynamic range data made possible by these devices.

References

[1] P. E. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH'97*, pages 369–378, 1997.

[2] M. D. Grossberg and S. K. Nayar. High Dynamic Range from Multiple Images: Which Exposures to Combine? In *Proc. ICCV Workshop on Color and Photometric Methods in Computer Vision (CPMCV)*, 2003.

[3] IEEE Computer Society. IEEE Standard for Binary Floating-Point Arithmetic. Number 754-1985. IEEE, 1985.

[4] J. R. Janesick. *Scientific Charge-Coupled Devices*. SPIE, 2001.

[5] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High Dynamic Range Video. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2003)*, 22(3):319–325, 2003.

[6] T. Lulé, H. Keller, M. Wagner, and M. Böhm. LARS II - A High Dynamic Range Image Sensor with a-Si:H Photo Conversion Layer. In *1999 IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, Nagano, Japan, 1999.

[7] B. C. Madden. Extended Intensity Range Imaging. Technical report, University of Pennsylvania, GRASP Laboratory, 1993.

[8] T. Mitsunaga and S. K. Nayar. Radiometric Self Calibration. In *Proc. CVPR-99*, pages 374–380. IEEE, 1999.

[9] M. A. Robertson, S. Borman, and R. L. Stevenson. Estimation-theoretic Approach to Dynamic Range Enhancement Using Multiple Exposures. *Journal of Electronic Imaging*, 12(2):219–228, April 2003.

[10] U. Seger, H.-G. Graf, and M. E. Landgraf. Vision Assistance in Scenes with Extreme Contrast. *IEEE Micro*, 12(1):50–56, 1993.

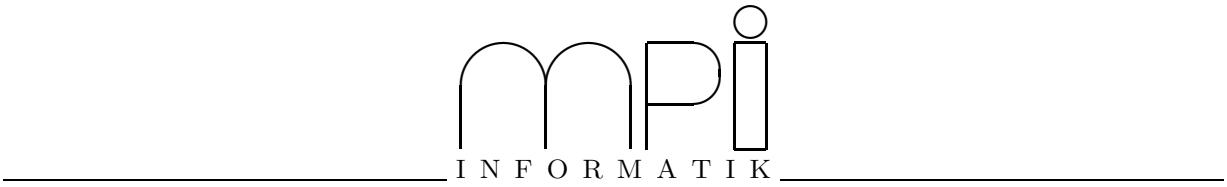

Below you find a list of the most recent technical reports of the Max-Planck-Institut für Informatik. They are available by anonymous ftp from `ftp.mpi-sb.mpg.de` under the directory `pub/papers/reports`. Most of the reports are also accessible via WWW using the URL `http://www.mpi-sb.mpg.de`. If you have any questions concerning ftp or WWW access, please contact `reports@mpi-sb.mpg.de`. Paper copies (which are not necessarily free of charge) can be ordered either by regular mail or by e-mail at the address below.

Max-Planck-Institut für Informatik
Library
attn. Anja Becker
Stuhlsatzenhausweg 85
66123 Saarbrücken
GERMANY
e-mail: `library@mpi-sb.mpg.de`

MPI-I-2005-4-006	C. Fuchs, M. Goesele, T. Chen, H. Seidel	An Emperical Model for Heterogeneous Translucent Objects
MPI-I-2005-4-005	G. Krawczyk, M. Goesele, H. Seidel	Photometric Calibration of High Dynamic Range Cameras
MPI-I-2005-4-004	C. Theobalt, N. Ahmed, E. De Aguiar, G. Ziegler, H. Lensch, M.A.,. Magnor, H. Seidel	Joint Motion and Reflectance Capture for Creating Relightable 3D Videos
MPI-I-2005-4-003	T. Langer, A.G. Belyaev, H. Seidel	Analysis and Design of Discrete Normals and Curvatures
MPI-I-2005-4-002	O. Schall, A. Belyaev, H. Seidel	Sparse Meshing of Uncertain and Noisy Surface Scattered Data
MPI-I-2005-4-001	M. Fuchs, V. Blanz, H. Lensch, H. Seidel	Reflectance from Images: A Model-Based Approach for Human Faces
MPI-I-2005-2-002	P. Maier, W. Charatonik, L. Georgieva	Bounded Model Checking of Pointer Programs
MPI-I-2005-2-001	J. Hoffmann, C. Gomes, B. Selman	Bottleneck Behavior in CNF Formulas
MPI-I-2005-1-007	I. Katriel, M. Kutz	A Faster Algorithm for Computing a Longest Common Increasing Subsequence
MPI-I-2005-1-002	I. Katriel, M. Kutz, M. Skutella	Reachability Substitutes for Planar Digraphs
MPI-I-2005-1-001	D. Michail	Rank-Maximal through Maximum Weight Matchings
MPI-I-2004-NWG3-001	M. Magnor	Axisymmetric Reconstruction and 3D Visualization of Bipolar Planetary Nebulae
MPI-I-2004-NWG1-001	B. Blanchet	Automatic Proof of Strong Secrecy for Security Protocols
MPI-I-2004-5-001	S. Siersdorfer, S. Sizov, G. Weikum	Goal-oriented Methods and Meta Methods for Document Classification and their Parameter Tuning
MPI-I-2004-4-006	K. Dmitriev, V. Havran, H. Seidel	Faster Ray Tracing with SIMD Shaft Culling
MPI-I-2004-4-005	I.P. Ivrissimtzis, W.-. Jeong, S. Lee, Y.a. Lee, H.-. Seidel	Neural Meshes: Surface Reconstruction with a Learning Algorithm
MPI-I-2004-4-004	R. Zayer, C. Rssl, H. Seidel	r-Adaptive Parameterization of Surfaces
MPI-I-2004-4-003	Y. Ohtake, A. Belyaev, H. Seidel	3D Scattered Data Interpolation and Approximation with Multilevel Compactly Supported RBFs
MPI-I-2004-4-002	Y. Ohtake, A. Belyaev, H. Seidel	Quadric-Based Mesh Reconstruction from Scattered Data
MPI-I-2004-4-001	J. Haber, C. Schmitt, M. Koster, H. Seidel	Modeling Hair using a Wisp Hair Model
MPI-I-2004-2-007	S. Wagner	Summaries for While Programs with Recursion

MPI-I-2004-2-002	P. Maier	Intuitionistic LTL and a New Characterization of Safety and Liveness
MPI-I-2004-2-001	H. de Nivelle, Y. Kazakov	Resolution Decision Procedures for the Guarded Fragment with Transitive Guards
MPI-I-2004-1-006	L.S. Chandran, N. Sivadasan	On the Hadwiger's Conjecture for Graph Products
MPI-I-2004-1-005	S. Schmitt, L. Fousse	A comparison of polynomial evaluation schemes
MPI-I-2004-1-004	N. Sivadasan, P. Sanders, M. Skutella	Online Scheduling with Bounded Migration
MPI-I-2004-1-003	I. Katriel	On Algorithms for Online Topological Ordering and Sorting
MPI-I-2004-1-002	P. Sanders, S. Pettie	A Simpler Linear Time 2/3 - epsilon Approximation for Maximum Weight Matching
MPI-I-2004-1-001	N. Beldiceanu, I. Katriel, S. Thiel	Filtering algorithms for the Same and UsedBy constraints
MPI-I-2003-NWG2-002	F. Eisenbrand	Fast integer programming in fixed dimension
MPI-I-2003-NWG2-001	L.S. Chandran, C.R. Subramanian	Girth and Treewidth
MPI-I-2003-4-009	N. Zakaria	FaceSketch: An Interface for Sketching and Coloring Cartoon Faces
MPI-I-2003-4-008	C. Roessl, I. Ivrissimtzis, H. Seidel	Tree-based triangle mesh connectivity encoding
MPI-I-2003-4-007	I. Ivrissimtzis, W. Jeong, H. Seidel	Neural Meshes: Statistical Learning Methods in Surface Reconstruction
MPI-I-2003-4-006	C. Roessl, F. Zeilfelder, G. Nrnberger, H. Seidel	Visualization of Volume Data with Quadratic Super Splines
MPI-I-2003-4-005	T. Hangelbroek, G. Nrnberger, C. Roessl, H.S. Seidel, F. Zeilfelder	The Dimension of C^1 Splines of Arbitrary Degree on a Tetrahedral Partition
MPI-I-2003-4-004	P. Bekaert, P. Slusallek, R. Cools, V. Havran, H. Seidel	A custom designed density estimation method for light transport
MPI-I-2003-4-003	R. Zayer, C. Roessl, H. Seidel	Convex Boundary Angle Based Flattening
MPI-I-2003-4-002	C. Theobalt, M. Li, M. Magnor, H. Seidel	A Flexible and Versatile Studio for Synchronized Multi-view Video Recording
MPI-I-2003-4-001	M. Tarini, H.P.A. Lensch, M. Goesele, H. Seidel	3D Acquisition of Mirroring Objects
MPI-I-2003-2-004	A. Podelski, A. Rybalchenko	Software Model Checking of Liveness Properties via Transition Invariants
MPI-I-2003-2-003	Y. Kazakov, H. de Nivelle	Subsumption of concepts in DL $\mathcal{FL}_0$ for (cyclic) terminologies with respect to descriptive semantics is PSPACE-complete
MPI-I-2003-2-002	M. Jaeger	A Representation Theorem and Applications to Measure Selection and Noninformative Priors
MPI-I-2003-2-001	P. Maier	Compositional Circular Assume-Guarantee Rules Cannot Be Sound And Complete
MPI-I-2003-1-018	G. Schaefer	A Note on the Smoothed Complexity of the Single-Source Shortest Path Problem
MPI-I-2003-1-017	G. Schfer, S. Leonardi	Cross-Monotonic Cost Sharing Methods for Connected Facility Location Games
MPI-I-2003-1-016	G. Schfer, N. Sivadasan	Topology Matters: Smoothed Competitive Analysis of Metrical Task Systems
MPI-I-2003-1-015	A. Kovcs	Sum-Multicoloring on Paths
MPI-I-2003-1-014	G. Schfer, L. Becchetti, S. Leonardi, A. Marchetti-Spaccamela, T. Vredeveld	Average Case and Smoothed Competitive Analysis of the Multi-Level Feedback Algorithm
MPI-I-2003-1-013	I. Katriel, S. Thiel	Fast Bound Consistency for the Global Cardinality Constraint
MPI-I-2003-1-012		- not published -
MPI-I-2003-1-011	P. Krysta, A. Czumaj, B. Voecking	Selfish Traffic Allocation for Server Farms

MPI-I-2003-1-010	H. Tamaki	A linear time heuristic for the branch-decomposition of planar graphs
MPI-I-2003-1-009	B. Csaba	On the Bollobás – Eldridge conjecture for bipartite graphs
MPI-I-2003-1-008	P. Sanders	Polynomial Time Algorithms for Network Information Flow