How Bright Is the Moon? Recovering and Using Absolute Luminance Values from Internet Images

Jens Ackermann and Michael Goesele

TU Darmstadt

Abstract. The human visual system differs from a camera in various aspects such as spatial resolution, brightness sensitivity, dynamic range, or color perception. Several of these effects depend on the absolute luminance distribution entering the eye which is not readily available from camera images. In this paper, we argue that absolute luminance is important for correct image reproduction. We investigate to which extent it is possible to recover absolute luminance values for any pixel in images taken from the Internet, extending previous studies on camera calibration in laboratory settings that are much less challenging. We use the Moon as a calibration target to estimate the remaining error. We then evaluate this error in the context of perceptual tonemapping for low dynamic range images.

Keywords: luminance, Internet photo collections, calibration.

1 Introduction

Today's consumer cameras produce visually pleasing images under standard scene conditions, i.e., images captured with common exposure characteristics. They work well as long as we are not interested in 'non-standard' scenes. E.g., images taken under low-light conditions will lead to long exposure times but produce a sharp, colorful image (disregarding any noise). This is, however, not the experience an observer would have in the actual scene and such images therefore tend to look artificial. To better reproduce human perception in these cases, one would need an estimate of absolute scene luminance. In addition, being able to recover absolute luminance values is also of interest in other areas such as computer vision. The motivating question for this paper is thus 'Can we estimate per-pixel luminance values for everyday Internet images?'.

In theory, the ISO standard 12232 [12] answers this question as it specifies a relation of camera output to scene exposure. Applying this model in practice, is however difficult since Internet images might be edited, contain false information, or because camera manufacturers do not follow the standard. But still, a lot of images should roughly comply with this model. The interesting question is therefore 'How well can we estimate per-pixel luminance values?'.

To answer this question, we need a large amount of images with ground truth luminance values. In a controlled environment, Wueller and Gabele [28] used a

luminance meter to measure the scene and then compare to a calibrated camera. But this does not match well with the variations in Internet images which are taken by many different camera models, show different scenes, or are possibly post-processed. Instead, our approach is to use images that were actually downloaded from platforms such as Flickr and Google images. Measuring the ground-truth with a luminance meter for these images might work for a frequently photographed, indoor scene without many luminance changes (e.g., an exhibit in a museum). But preferably, we would like to even dispense with the need of measuring the ground truth and instead use a calibration target with known luminance.

We found such a target in the Moon and assume that insights gained on lunar images from the Internet transfer to the more general setting. Working with lunar images and different models of lunar photometry, we discovered, however, that observations of the Moon from Earth are not as predictable as expected. Atmospheric effects and unknown observer positions can lead to a significant variation that is beyond our control (in addition to the uncertainty in the camera model). To demonstrate the applicability of the approach to a practical problem we implement a re-mapping approach of low dynamic range images that transforms artificially looking low-light pictures into a perceptually plausible representation.

2 Related Work

Camera Calibration: Radiometric camera calibration is important for many computer vision tasks that rely on linear luminance measurements, e.g., high dynamic range reconstruction, image stitching, or photometric stereo. Accordingly, there is a large body of literature on this topic, e.g., [22,9,18,19,15]. These papers focus mostly on the calibration up to relative luminance values and typically consider a single camera.

Martinez-Verdu et al. [25] perform a series of measurements to calibrate a digital camera. After their colorimetric characterization a standard camera can be used as absolute colorimeter. They discover that predicted colors from the camera differ from the true values according to an affine model. The characterization does of course not scale to the multitude of cameras that are used to take pictures on the Internet. We abandon these rigorous measurements and replace them with a less reliable but more general model given by the ISO standard for camera sensitivity. Wueller and Gabele [28] calibrate the opto-electronic conversion function of several digital cameras for a certain exposure level. They can then transform measurements under different exposure settings to the calibrated, absolute luminance values. Comparing these values in different scenes to those measured with a luminance meter they find possible deviations of more than 30% for colored objects. We must assume that these deviations will even increase in the case of uncontrolled Internet images.

Recently, the need to make predictions of radiometric calibration on Internet data has arisen since Internet Photo Collections have become an input source for computer vision algorithms. Chakrabarti et al. [3] study the imaging pipeline

from scene luminance to final pixel value in this context. They propose a camera model that also encompasses the scene-dependent variation, but consider only relative luminance values. Kuthirummal et al. [17] analyze the statistics of images in photo collections grouped by camera model and lens settings. They derive a prior on the joint histogram of irradiances at neighboring pixels and can then estimate response functions for other camera models in the set. Again, this only recovers luminance values up to an unknown scaling factor.

The ROLO (RObotic Lunar Observatory) program [1] by USGS is aimed at providing radiometric calibration for space-based imaging instruments, using the Moon as a reference source. We follow this idea but apply it in a totally different context: We assume Earth-based consumer cameras instead of specialized space equipment and do not know the actual time of capture or the observer’s location.

Luminance of the Moon: The apparent brightness of the Moon has been of interest for centuries since it is the second brightest object in the sky after the sun. We will not discuss this in detail and mention only selected works. Ellis [6] combines previous results about the relative intensity of the Moon with atmospheric extinction to derive tabulated intensities for varying zenith angle and lunar phase. These values are relative to the intensity of the full moon and cannot be used directly for absolute calibration.

In 1994, the Clementine mission produced high resolution images of the lunar surface from a polar orbit. This enabled an updated model of the lunar albedo at various wavelengths and a study of the so called *opposition effect* by Buratti et al. [2]. The opposition effect leads to an apparent increase of albedo for small phase angles which has also been studied by Hapke [10]. Using Hapke’s work, Jensen et al. [13] assembled a model of the complete night sky including the Moon which can be used for realistic image synthesis. Their formulation is centered on irradiance which we transform into radiance/luminance at a single camera pixel. This is the inverse of Kieffer and Stone’s approach [14] processing lunar radiance measurements in the context of the above mentioned ROLO program.

Finally, in working with Internet images we also encountered pictures of lunar eclipses since these draw special attention. Hernitschek et al. [11] observe eclipses with a light meter and devise a model for the luminance of the Moon while in the Earth’s shadow. Since this is a special case, we do not use images of lunar eclipses and consider them as outliers in our derivations.

Perceptual Tonemapping: As an application of absolute luminance in images we study perceptual effects in human vision that depend on the absolute adaption state of the eye’s receptor cells. Spencer et al. [24] study glare effects arising from scattering in ocular media and diffraction at the iris. Ward et al. [27] as well as Ferwerda et al. [7] model glare, color sensitivity, and visual acuity in the context of rendered high dynamic range scenes. Their focus is on employing local contrast thresholds for dynamic range compression which also depends on absolute luminance. Durand and Dorsey [5] extend on this work but still consider rendered images. They use a global adaption level per image and compare its estimation to a camera’s exposure metering.

With the availability of high dynamic range photographs through multiple exposure algorithms, tone-mapping has also been applied to actual photographs. Reinhard et al. [21] use relative luminance values and a *key value* that adjusts the overall brightness. They do not explicitly model any perceptual effects. These are then introduced to photographic tone-mapping of high dynamic range videos by Krawczyk et al. [16]. We show that effects such as reduced acuity in low-light conditions can also be employed with normal low dynamic range images to increase their ‘naturalness’.

3 Camera Model

We model a camera as consisting of a linear part defined by its optical system and linear sensor, and a non-linear processing part that may contain everything from sensor electronics to post-processing or image editing. The optical system transforms an incoming scene luminance L into the focal plane exposure H which is proportional to the integration time t , but decreases with the f-number N of the lens. The proportionality factor according to annex B of the ISO standard 12232 [12] yields:

$$H = 0.65 \frac{L \cdot t}{N^2} . \quad (1)$$

Fortunately, a lot of images contain the values of t and N in their EXIF meta-data. The signal is then scaled by the sensor gain and processed by the camera to yield non-linear 8-bit pixel values $p = 255 \cdot f(\tilde{p})$. We assume f to be a gamma curve as defined by the sRGB standard which is commonly used for Internet images. The missing link between a certain absolute exposure H and an output pixel value p is then just the sensitivity of the sensor.

ISO standard 12232 [12] includes several definitions of sensitivity. But camera vendors do not always comply with this standard exactly. Even if they do, it is usually not apparent from the available meta-data which definition was used. This can lead to greatly differing exposures for the same pixel value. A value of 200 for the ISO speed interpreted as *standard output sensitivity* implies that $H = 10/200$ gets mapped to a pixel value of 118. If it is interpreted instead as *saturation-based speed*, a pixel value 118 corresponds to an exposure of $f^{-1}(118/255) \cdot 78/200 = 0.1835 \cdot 78/200$ which is 1.43 times the previous one.

According to the EXIF standard, the sensitivity type should be recorded along with the sensitivity value. Unfortunately, many images contain only the ‘ISO’ tag and do not specify which definition it relates to. We assume that the standard output sensitivity (*SOS*) based speed S as defined in Sect. 7.1 of the standard can be applied for all images. Under these assumptions the luminance L_{SOS} that produces an output value of $255 \cdot f(\tilde{p}) = 118$ is given as

$$L_{SOS} = \frac{1000}{65} \cdot \frac{N^2}{S \cdot t} \quad (2)$$

More generally, a pixel value p corresponds to luminance

$$L = \frac{f^{-1}(p/255)}{f^{-1}(118/255)} \cdot L_{SOS} = \tilde{p} \cdot 5.45 \cdot \frac{1000}{65} \cdot \frac{N^2}{S \cdot t} . \quad (3)$$

Since we typically deal with color images, we cannot use this formula directly. Instead, we compute grayscale values by combining the RGB channels according to Sect. 6.3.3 in ISO 12232:

$$\tilde{p} = 0.2125 \cdot f^{-1} \left(\frac{R}{255} \right) + 0.7154 \cdot f^{-1} \left(\frac{G}{255} \right) + 0.0721 \cdot f^{-1} \left(\frac{B}{255} \right). \quad (4)$$

4 The Moon as Calibration Target

We make use of the lunar illumination model in [13] that relates solar irradiance on the Moon $\tilde{E}_{sm}$ to the irradiance on Earth arriving from the Moon $\tilde{E}_m$ depending on the phase angle φ :

$$\tilde{E}_m(\varphi) = \frac{2}{3} \cdot \frac{C}{\pi} \cdot \omega_M \cdot \tilde{E}_{sm} \cdot (1 - \sin(\varphi/2) \tan(\varphi/2) \log(\cot(\varphi/4))) \quad (5)$$

We ignore the earthshine component which is only relevant around a new moon. The average albedo C is set to 0.072 and ω_M is the solid angle of the complete lunar disk as observed from Earth. From our camera model we get luminance values instead of irradiances. Exchanging radiometric quantities with photometric ones in (5) gives us the illuminance E_m . The illuminance integrates incoming luminance L over the whole hemisphere:

$$E = \int L(\theta) \cos(\theta) d\omega \quad (6)$$

Assuming that L is constant for all directions θ which correspond to the illuminated portion of the Moon Ω_P , zero everywhere else, and that the sensor is oriented towards the Moon we get

$$E_m(\varphi) = L \int_{\Omega_P} \cos(\theta) d\omega \approx L \int_{\Omega_P} d\omega = L \cdot \omega_P. \quad (7)$$

The solid angle ω_P can be computed from the average radius $r_M = 1737 \text{ km}$ of the Moon and its average distance from Earth $d = 384400 \text{ km}$

$$\omega_P = P \cdot \frac{\pi r_M^2}{d^2} = P \cdot \omega_M \quad (8)$$

where P is the percentage of the lunar disc that is illuminated as seen from Earth. P depends on the phase angle as follows:

$$P = 0.5 \cdot (1 + \cos(\varphi)). \quad (9)$$

Thus, our final model of extra-terrestrial, lunar luminance is

$$L_m = \frac{E_m}{\omega_P} = \frac{2}{3} \cdot \frac{C}{\pi} \cdot \frac{2 E_{sm}}{1 + \cos(\varphi)} \cdot (1 - \sin(\varphi/2) \tan(\varphi/2) \log(\cot(\varphi/4))). \quad (10)$$

Fig. 1. *Left:* Extra-terrestrial luminance for various phase angles φ . *Right:* Expected luminance distribution at sea level after atmospheric extinction for $\varphi = 5^\circ, 25^\circ, 90^\circ$.

We use a value of $1.338 \cdot 10^5 \text{ lm/m}^2$ for the illuminance of the sun E_{sm} at a distance of one astronomical unit ($1 \text{ AU} = 149597870.7 \text{ km}$) as in the CIE sky model [4]. Ignoring the minor deviations of the actual position of the Moon from 1 AU Fig. 1 shows the expected luminance values for a range of phase angles. Note that we model luminance and not brightness or irradiance. The latter would fall off much faster, because the illuminated area of the Moon decreases.

So far, we have only considered luminances as would be observed in free space, but the luminance gets attenuated by the atmosphere depending on the amount of air between the observer and the Moon. At sea level, the extinction due to Rayleigh scattering R_{Ray} and aerosol scattering R_{aer} can be approximated according to Green [8] as

$$F = 2.512^{R_{Ray}+R_{aer}} \cdot A = 2.512^{0.1451+0.12} \cdot A = 1.2766 A. \quad (11)$$

The air mass A is defined as 1 when looking straight up (zenith) and can be calculated from the zenith angle ϕ of the Moon at the observer position:

$$A = (\cos(\phi) + 0.025 \exp(-11 \cos(\phi)))^{-1}. \quad (12)$$

Ignoring any wavelength dependent effects, the final observed luminance will be $L = L_m/F$.

Since we do not know where an image was taken, we cannot deduce the zenith angle or the approximate extinction coefficient. To get an impression of the severity of this error source, we assume an exemplary Gaussian distribution of zenith angles centered at 55° with a standard deviation of 25° . This reflects that the Moon is seldom observed at the zenith, but more often close to the horizon. Fig. 1 then shows the expected luminance distributions for phase angles $\varphi = 5^\circ, 25^\circ, 90^\circ$. For each phase curve, luminances can vary almost over the complete range indicating that atmospheric extinction is a major source of error.

5 Predicting Luminance from Internet Images

We downloaded an initial set of about 13000 images from Flickr with tags ‘moon’, ‘full’, ‘night’, and ‘sky’. Those without the necessary EXIF data were automatically removed. Many images show artistic works or landscape shots with the

Fig. 2. *Left:* Histogram (blue) of predicted $\log_{10}$ luminances from image meta data with fitted gaussian (red, $\mu = 3.25, \sigma = 0.4$). *Right:* Histogram of differences of both predictions ($\mu = 0.04, \sigma = 0.4$).

Fig. 3. *Left to right:* The geometry of lunar illumination coverage, an example image of the Moon, the binary mask after thresholding, initial coverage (red) assumed as initialization for the optimization (image mask overlaid in white), estimated coverage after optimization ($\varphi = 70^\circ$)

Moon as a small, over-exposed disk. We removed most of these by comparison with an actual photograph of the Moon using the PictureRelate tool [26]. The remaining 800 images were uploaded to Amazon Mechanical Turk where humans annotated them with a bounding box of the Moon.

Each image was then transformed according to (4) and thresholded at $\tilde{p} = 0.3$ to create a binary mask of the Moon. We need the phase angle to evaluate the model in (10) and predict luminances in each image. This angle can be estimated from a geometric model of the illumination coverage of the lunar disk. The model is illustrated in Fig. 3 and accounts for unknown pixel coordinates of the center (x, y) , the radius, a tilt angle, and the phase. Its output is the set of pixels that correspond to the illuminated part of the Moon (red in Fig. 3). We run an optimization that fits this model to the binary mask in each image and returns the estimated lunar phase angles. We can then evaluate (10) to predict the extra-terrestrial luminance values L_{lunar} . These are in the range $[1360, 2040] \text{ cd/m}^2$ as indicated by Fig. 1. For the reference image shown in Fig. 3 we know the exact date and position of capture. Comparing the estimated phase angle of $\varphi = 70^\circ$ with data provided by NASA’s HORIZONS system [20] $\varphi = 65^\circ$ yields a deviation that is negligible for our derivations.

Applying (3), we can also compute luminance estimates based on the actual pixel values and meta-data. We first exclude pixels with $\tilde{p} \geq 250/255$ as over-

exposed and then compute per-pixel luminance estimates within the mask. These are then averaged to yield the predicted luminance of the Moon based on image measurements L_{image} . Note that these are Earth-bound and may include atmospheric effects. The histogram of image-based predictions greater than 30 cd/m^2 is displayed in Fig. 2 (left). We also removed images with more than 1% overexposed pixels which leads to a total of 395 images.

Finally, we compare both predictions by taking the ratio L_{lunar}/L_{image} . The respective histogram is plotted in Fig. 2 (right). We observe that the mean of both predictions agrees well whereas the standard deviation yields a factor of $10^{-0.4} \approx 0.4$ or $10^{0.4} \approx 2.5$, respectively. Note that this number includes all variations and unmodeled effects in the complete pipeline. Common examples are images with cloudy or not completely black sky, e.g., during dusk.

6 Perceptual Remapping of Internet Images

In the previous sections we have seen that it is possible to reconstruct absolute luminance values for each pixel in an image from meta-data. Absolute luminances are important for a lot of perceptual effects that depend on the absolute adaption level. Cameras aim at reproducing images an observer has seen in photopic conditions. But important changes in perception occur in the mesopic and scotopic range of low-light illumination. A camera will therefore fail to produce plausible images under such conditions. We will focus on two such effects and show how absolute luminances can be employed to re-render night images in a perceptually more plausible way.

Visual Acuity: Acuity expresses the spatial resolution of the visual system. This can also be seen as sharpness and is typically measured in cycles per visual degree. For human vision, a basic test from ophthalmology is for example to have the patient read letters of decreasing size at a fixed distance. It has been shown by Shaler [23] that performance in these tests depends on the adapting luminance. Under moonlight conditions, acuity drops significantly and spatial detail is lost even for healthy observers.

Ward et al. [27] derive the following relation of luminance to maximal resolvable frequency from Shaler’s data:

$$f(L) = 25.72 + 17.25 \cdot \tan^{-1} (1.4 \log_{10}(L) + 0.35). \quad (13)$$

To simulate this effect we blur the image with a low-pass filter that removes frequencies above $f(L)$. Note that the adaption luminance is not constant over the whole image, but may change locally. Different levels of blur will be appropriate in different regions of the image. We borrow from Krawczyk et al. [16] and use the same kernels as in their high dynamic range work.

Rod and Cone Activity: Another effect that becomes apparent in low-light is the loss of color. In the scotopic range of about 10^{-6} to 10^{-2} cd/m^2 the cone cells in the human eye are inactive and vision is solely based on the rod photoreceptor

Fig. 4. Stages of perceptual simulation. *Top:* The original image taken with a long exposure time shows bright colors and sharp details (left). We remove spatial frequencies that could not be resolved by an observer under low-light conditions (right). *Bottom:* Rod responses are simulated through a grayscale image (left) which is then blended with the colored one according to local rod activity. The final result (right) shows both effects combined.

cells. Again, cameras mimic the human eye under photopic conditions when the three types of cones would provide a perception of color.

The rods have a spectral sensitivity that differs from the luminosity function of the cones which is usually applied in photometry. Without knowing the spectral distribution of incoming light, it is not possible to directly translate one into the other. We therefore neglect this effect, use photopic luminance values in both cases, and implement the loss of color perception by blending a grayscale image and a colored one. The blending coefficient simulates the decrease in rod activity with increasing adaption luminance. Again, these coefficients can vary spatially over the image.

Ferwerda et al. [7] propose a linear blending weight for the mesopic range whereas Durand and Dorsey [5] employ the ratio of two linear functions of adaption luminance. We use the same method as Krawczyk et al. [16] who define the blending weight as

$$\sigma(L) = \frac{0.04}{0.04 + L}, \quad (14)$$

effectively disabling rod contribution for luminances above 10 cd/m^2 .

Fig. 5. Results. *Left:* The original, unprocessed images (rows 1-3: Flickr users Joselito Tagarao, Markus Lehr, Wayne Grivell). *Right:* Results show-casing reduced acuity and color perception in low-light conditions. The daylight picture in Fig. 6 is not affected.

Implementation: We assume an sRGB image as input and first transform it into the xyY space. Using (3) we scale the Y channel to obtain absolute luminance values. We then build a Gaussian pyramid of successively blurred luminance images similar to Krawczyk et al. [16] but without their approximation of kernels since we are not concerned about real-time performance. Based on the absolute luminance and (13) we select the appropriate pyramid level for each pixel and obtain a blurred luminance map L_{ac} .

To store the result, we need to arrive at relative pixel values. Reversing (3) gives $\tilde{L}_{ac}$ which is then split into a scotopic part $Y_s = \sigma(L_{ac}) \cdot \tilde{L}_{ac}$ and a photopic one $Y_p = (1 - \sigma(L_{ac})) \cdot \tilde{L}_{ac}$. The blending coefficient still depends on absolute luminances. Finally, we transform the original chromaticities xy with Y_p into the XYZ space and add the grayscale image obtained from Y_s and $x = 1/3, y = 1/3$.

Fig. 4 shows the different images occurring in this algorithm. The original was taken in a dark capture lab with the white glove still recognizable by a human. The color checker became visible after a short period of dark adaption. The sharp and colored camera picture with an exposure time of 15 s is very far from the experience a human would have in this setting.

Fig. 6. Results on daylight image. *Left:* The original image. *Right:* Results of applying the pipeline. As expected, the daylight picture is not affected.

Fig. 7. Colormapped $\log_{10}$ of absolute luminance. Images from left to right correspond to rows in Fig. 5 and Fig. 6. The respective logarithmic mean of estimated luminances is: 0.0003 cd/m^2 , 0.014 cd/m^2 , 0.018 cd/m^2 , 92.86 cd/m^2 . The estimated luminance factor is: 0.087 cd/m^2 , 0.576 cd/m^2 , 0.407 cd/m^2 , 1128.23 cd/m^2 .

Fig. 8. Impact of possible false estimates. *Left to right:* results for an additional scaling factor of 0.2, 0.5, 1, 2, 5 in the luminance estimation.

Results on Internet Images: To test our simulation of human perception, we downloaded images from Flickr that were taken with only little light, but still look sharp and colorful due to long exposure times. It is immediately apparent that the original images in the left column of Fig. 5 look artificial. This is of course intended and an artistic choice by the photographer, but we are interested in how a human would actually have perceived the scene. In our results in the right column, the loss of acuity can be well observed on the surface of the rock in the first row, at the air vent in the second row, and at the fence in the third row. Similarly, colors are much less saturated. Note especially the green plants and tufts of grass in the third row and the overall appearance in the first row. In general, our images look more plausible with regard to a night time scene than the original long-exposure shots. Of course, day light images should not be affected by our simulation and we show one such example in Fig. 6.

In these examples, we are dealing with a range of luminances from starlight at about 0.001 cd/m^2 to clouded sky illumination of 1000 cd/m^2 . Fig. 7 shows false color images of the absolute log luminances in all four images. To study the impact that possible false estimates of the luminance factor could have on overall image appearance, we artificially introduced an additional scale factor α which we varied among 0.2, 0.5, 1.0, 2.0, 5.0. Fig. 8 shows that deviations by a factor of 5 are clearly discernible. For $\alpha = 2$ we can make out a difference but the overall impression of a low-light scene is still preserved.

7 Discussion

We originally posed the question of how well we can estimate absolute luminance values on uncontrolled, real-world images. Using images of the Moon downloaded from the Internet we compute luminance values once through a photometric model of the Moon and again through an image formation model based on the available information. Ideally, both estimates would agree at least to the uncertainty within each of these models. We can, however, not control or even estimate several parameters of these models such as atmospheric extinction, the way of reporting ISO settings, or other effects that are “invisible” (e.g., filters in front of the lens, photographing through a telescope). This makes it impossible to exactly quantify the uncertainties in the overall system. We are also aware of the fact that our camera model is an inexact approximation in many cases. The major concerns are incorrect meta-data, post-processed images, non-standard response curves, and vendor-specific changes. We briefly tested the latter two factors by capturing a color checker with our Canon EOS camera and simultaneously acquiring measurements with a luminance meter. It turned out that our predicted luminances according to Sect. 3 deviated approximately $\pm 30\%$ from ground truth. This matches with the findings by Wueller and Gabele [28] who even calibrated the camera’s response.

Can we still gain insights from our investigations? Yes, since the discrepancies we observe, e.g., illustrated in Fig. 2, are an upper bound for the question of how accurately we can recover absolute luminances. If the deviations stemmed purely from the photometric model, then luminance recovery from meta-data would be precise. If on the other hand the photometric model was precise, the observed variance would be the answer to our initial question. The truth will be somewhere between these extrema, but our investigations still provide an upper bound.

Finally, we can refine our motivating question and ask ‘How accurate do we have to be?’ This of course depends on the application. In this work, we have studied the case of perceptual re-mapping of low-light images which clearly improves the realism of very dim images but becomes only possible through our estimated luminances. In this context, a factor of 2 (or one exposure stop in photography terms) seems to be still acceptable (see Fig. 8). Affirming this impression through a user study would be an interesting avenue for future work.

Acknowledgments. This work was supported in part by the DFG Emmy Noether fellowship GO 1752/3-1. The authors want to thank Carsten Haubold and Fabian Langguth for an early analysis of the lunar images. Philipp Urban and Jens Preiss shared their experience on camera calibration. Karol Myszkowski and Grzegorz Krawczyk provided valuable discussions and parts of their source code. We also thank all Flickr users who contributed to this project.

References

1. Lunar calibration, USGS robotic lunar observatory (ROLO), www.moon-cal.org
2. Buratti, B., Hillier, J., Wang, M.: The lunar opposition surge: Observations by Clementine. *Icarus* (1996)
3. Chakrabarti, A., Scharstein, D., Zickler, T.: An empirical camera model for internet color vision. In: *BMVC*, pp. 51.1–51.11 (2009)
4. Darula, S., Kittler, R.: CIE general sky standard defining luminance distributions. In: *Proc. eSim* (2002)
5. Durand, F., Dorsey, J.: Interactive tone mapping. In: *EGWR* (2000)
6. Ellis, D.: Illumination Received from the Moon. *Journal of the Royal Astronomical Society of Canada* (1966)
7. Ferwerda, J.A., Pattanaik, S.N., Shirley, P., Greenberg, D.P.: A model of visual adaptation for realistic image synthesis. In: *SIGGRAPH*, pp. 249–258 (1996)
8. Green, D.W.E.: Correcting for atmospheric extinction. *International Comet Quarterly* (1992)
9. Grossberg, M., Nayar, S.: Modeling the space of camera response functions. *PAMI* 26(10), 1272–1282 (2004)
10. Hapke, B.: An Improved Theoretical Lunar Photometric Function. *The Astronomical Journal* (1966)
11. Hernitschek, N., Schmidt, E., Vollmer, M.: Lunar eclipse photometry: absolute luminance measurements and modeling. *Applied Optics* 47(34), 62–71 (2008)
12. ISO: Photography - digital still cameras - determination of exposure index, iso speed ratings, standard output sensitivity, and recommended exposure index (2006)
13. Jensen, H.W., Durand, F., Stark, M.M., Shirley, P., Dorsey, J., Premoze, S.: A physically-based night sky model. In: *SIGGRAPH* (2001)
14. Kieffer, H.H., Stone, T.C.: The spectral irradiance of the moon. *The Astronomical Journal* 129(6), 2887–2901 (2005)
15. Kim, S.J., Pollefeys, M.: Robust radiometric calibration and vignetting correction. *PAMI* 30(4), 562–576 (2008)
16. Krawczyk, G., Myszkowski, K., Seidel, H.P.: Perceptual effects in real-time tone mapping. In: *SCCG* (2005)
17. Kuthirummal, S., Agarwala, A., Goldman, D.B., Nayar, S.K.: Priors for Large Photo Collections and What They Reveal about Cameras. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008, Part IV. LNCS*, vol. 5305, pp. 74–87. Springer, Heidelberg (2008)
18. Litvinov, A., Schechner, Y.Y.: Addressing radiometric nonidealities: A unified framework. In: *CVPR*, pp. 52–59 (2005)
19. Matsushita, Y., Lin, S.: Radiometric calibration from noise distributions. In: *CVPR* (2007)
20. NASA, <http://ssd.jpl.nasa.gov/?horizons>
21. Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.A.: Photographic tone reproduction for digital images. *ACM Transactions on Graphics* (2002)
22. Robertson, M.A., Borman, S., Stevenson, R.L.: Dynamic range improvement through multiple exposures. In: *ICIP*, pp. 159–163 (1999)
23. Shaler, S.: The relation between visual acuity and illumination. *Journal of General Psychology* (1937)
24. Spencer, G., Shirley, P., Zimmerman, K., Greenberg, D.P.: Physically-based glare effects for digital images. In: *SIGGRAPH*, pp. 325–334 (1995)

25. Verdú, F.M., Pujol, J., Vilaseca, M., Capilla, P.: Characterization of a digital camera as an absolute tristimulus colorimeter. *J. of Imaging Science and Technology* 47(4), 279–295 (2003)
26. Walthelm, A.: <http://www.walthelm.net/picture-relate/index.php>
27. Ward Larson, G., Rushmeier, H., Piatko, C.: A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE TVCG* 3(4) (1997)
28. Wüller, D., Gabele, H.: The usage of digital cameras as luminance meters. In: *Proc. SPIE*, vol. 6502, pp. 1–11 (2007)