

Sky Based Light Metering for High Dynamic Range Images

Yulia Gryaditskaya¹, Tania Pouli^{1,2}, Erik Reinhard^{1,2} and Hans-Peter Seidel¹

¹Max-Planck Institute for Informatics, Saarbrücken, Germany

²Technicolor Research and Innovation, Rennes, France

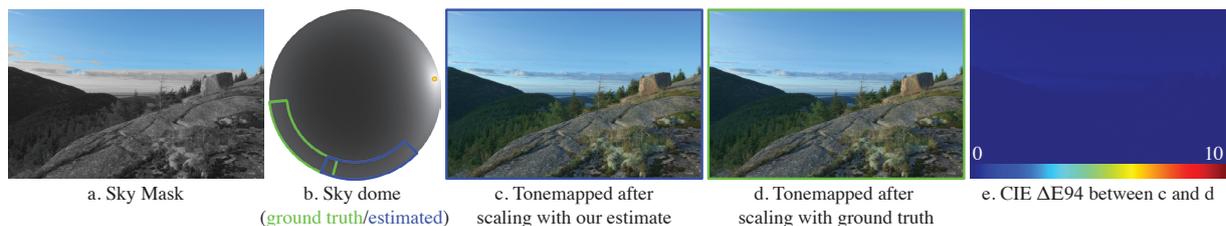


Figure 1: Our algorithm employs sky models to estimate a scaling that can convert images to absolute luminance values. (a) We detect and segment sky pixels, which are (b) analysed to determine properties of the sky dome and camera parameters (the boxes show the visible area of the sky dome in the image). (c) By scaling the image with our estimation, we obtain absolute radiometric values, allowing images to be processed by appearance models such as Reinhard et al. [RPK* 12]. (d) The same processing was applied to the image after scaling it with the ground truth scale factor. (e) The CIE ΔE_{94} differences between (c) and (d) indicate good correspondence between our estimate and the ground truth.

Abstract

Image calibration requires both linearization of pixel values and scaling so that values in the image correspond to real-world luminances. In this paper we focus on the latter and rather than rely on camera characterization, we calibrate images by analysing their content and metadata, obviating the need for expensive measuring devices or modeling of lens and camera combinations. Our analysis correlates sky pixel values to luminances that would be expected based on geographical metadata. Combined with high dynamic range (HDR) imaging, which gives us linear pixel data, our algorithm allows us to find absolute luminance values for each pixel—effectively turning digital cameras into absolute light meters. To validate our algorithm we have collected and annotated a calibrated set of HDR images and compared our estimation with several other approaches, showing that our approach is able to more accurately recover absolute luminance. We discuss various applications and demonstrate the utility of our method in the context of calibrated color appearance reproduction and lighting design.

Categories and Subject Descriptors (according to ACM CCS): I.3.m [Computer Graphics]: Image Processing—Computational photography

1. Introduction

Image processing and computer vision algorithms often require their input images to have a linear relationship between pixel values and scene luminance. Although knowledge of the camera curve is sufficient for creating such linear content, linearized images can only represent scene luminance up to a constant. Recently, however, several developments

have led to an increasing need for image content that is not only linear, but also absolute. This occurs whenever images need to be accurately reproduced on a high dynamic range display, when the image-based recovery of surface reflection parameters is performed or when material characteristics are measured, as well as when interactions between illumination and materials need to be accurately simulated.



Figure 2: Example renderings using skydomes that were scaled using our proposed algorithm.

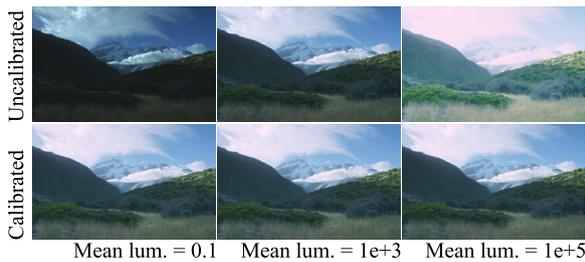


Figure 3: An HDR image was scaled to different mean luminances and processed using an appearance model [RPK* 12]. The top images were not calibrated, while the bottom images were scaled using our estimation. Despite the differences in mean luminance, the absolute scaling estimated using our algorithm remains consistent. The additional results are present in supplemental materials.

Additionally, in image-based lighting [Deb02], light probes captured from real scenes are used to realistically illuminate artificial scenes. Although such light probes are linear, they are not normally specified in absolute units. With absolute light probes, however, image-based lighting could become a powerful tool for lighting designers or architects (Fig. 2). Absolute values are required when one is interested in materials behaviour under different lighting conditions. The use of light probes would be beneficial over rendered skies [PSS99, HW12] due to variety of sky conditions, which cannot be reproduced by analytical sky-models.

Similarly, absolute scene luminance would be required in image reproduction, and in particular color appearance modelling. Here, the scene represented by the image would induce a specific state of adaptation in the observer. The display and its environment would induce a different state of adaptation. Color appearance models account for this mismatch in adaptation [Fai05] but we can only use such algorithms effectively if we know absolute light levels (Fig. 3).

Further, we consider the development of HDR technologies on mobile devices. A growing number of approaches

allows object insertion and relighting from a single image [KHFH11, MN08]. These could benefit from the input to be scaled to absolute values.

To calibrate images to accurate radiometric values, they need to be linearized first, and then scaled. In our case, we use high dynamic range (HDR) images, which are created from multiple exposures. As the camera response curve is estimated and accounted for during this process, such HDR images are linear by construction.

Although linearization of images has received considerable attention, algorithms for estimating absolute luminance values in images are few and far between, in particular those that do not require extraneous equipment. Conventionally, to obtain precise absolute values from the image, an 18% grey card or a color checker is inserted into the scene which is measured by a photometer. While this approach can be used in professional settings when capturing new imagery, the equipment necessary is not universally available and the measurements cannot be taken in retrospect. Consequently, when dealing with existing content (such as light probe galleries used for IBL), only the information within the image and its metadata is available. Moreover, in conditions of fast varying outdoor illumination (clouds, sunset, sunrise, storm etc.) when the photographer is interested in specific lightening phenomena, measurements of the scene are likely to be unreliable or difficult to make.

Another strategy to get the images in absolute values is to perform an initial camera calibration [WG07, FH09], which is valid only for a particular combination of a camera and a lens. In order to perform the initial calibration the approach still requires measurement equipment. Moreover, these techniques can not be applied to HDR images, as they estimate absolute values per pixel for the LDR image based on each exposure value. Without measurements, the absolute values can be obtained under several assumptions about a scene and a camera using EXIF based approaches, which are discussed in [AG13]. We discuss how these approaches can be applied to HDR images, as well as show the cases when these approaches can not be reliable and another solution is required.

In this paper, we explore the possibility of recovering absolute luminance values for outdoor images without the need for professional measurement hardware, relying instead on the camera itself as well as a GPS module, which is often included in new camera models.

To validate our algorithm, we require a set of calibrated HDR images with associated EXIF and GPS metadata. To this end, we have acquired a set of ground-truth images, which we aim to make generally available. In summary, our work offers the following contributions and benefits:

- No specialist equipment is required, nor is there a need to characterize the camera or lens.
- No EXIF or GPS data is required for environment maps with a fully visible sky dome.

- Our algorithm enables the recovery of absolute values with accuracy comparable to what could be obtained from direct scene measurements.
- A calibrated and annotated HDR dataset of outdoor scenes, which includes GPS and EXIF metadata, will be made publicly available.

2. Algorithm

The goal of our algorithm is to convert a linear HDR image of an outdoor scene into absolute radiometric measurements. To achieve that we compute a scale factor ξ , which converts relative luminance values to absolute luminance. The input to our algorithm is a single linear HDR image, GPS data and a small set of metadata, which can be extracted from the EXIF data of the image itself (date and time of capture, sensor size, focal length). If an environment map is used no GPS or EXIF information is required. Figure 4 illustrates the flow of the algorithm.

To compute ξ , we rely on sky regions in the image, found using a semantic image segmentation algorithm [HEH07]. Unlike most parts of an outdoor scene, the appearance of the sky is predictable and can be modelled with reasonable accuracy given a small set of parameters. This observation forms the main motivation for our framework; since pixels within an HDR image are linearly related, it is sufficient to estimate a scaling factor for sky pixels in the image.

Our algorithm first estimates the zenith luminance l_z of the sky dome relative to the sky pixels visible in the image. This is obtained by fitting the intensity values of the observed pixels to a parametric sky model [PSM93] and using the zenith value as a normalization factor to the range of the values in the image. During the fitting procedure the unknown parameters for the sky model are estimated as well as l_z . Using information obtained during this step, as well as some of the information within the metadata, we can also obtain an absolute estimate of the zenith luminance L_z . Since the original zenith estimate is given relative to the values of sky pixels in the image, this second absolute estimate allows us to compute a scale factor, that can finally be applied to all pixels:

$$\xi \triangleq L/l = L_z/l_z, \quad (1)$$

where L denotes absolute and l relative luminance values.

2.1. The fitting problem

The key idea of our approach is to extract information from the image by fitting a parametric sky model to the intensity values of the observed sky region. In previous works, sky regions were used to estimate camera spectral sensitivity and white balancing [KZTI13]. Using time-lapse sequences, skies were also used as geometric calibration targets to reverse-engineer intrinsic camera and lens parameters, including the focal length of the lens, the orientation of

$\theta_{c,s}$	Zenith angle (camera, sun)
ϕ_c	Camera azimuth angle
(u_p, v_p)	Coordinates of sky pixel in image
t	Turbidity
\vec{p}	Perez sky model parameters

Table 1: Symbols used in the description of the algorithm.

the camera as well as its geolocation [LGE10]. In a sense, our work is complementary to this analysis.

The appearance of the sky is caused by sunlight scattering through interactions with particles and molecules. Thus, sky luminance distribution models depend on the sun position as well as parameters which model the appearance of the sky under different atmospheric conditions. In our work, we use the Perez sky model as a prior [PSM93], which was shown to out-perform competing algorithms [IM94]. In the following, this model is represented by function $f(\cdot)$. Function $g(\cdot)$ restates the same model in terms of camera parameters [LGE10]. The fitting problem can then be formulated as an optimization (see Table 1):

$$\min_{\theta_c, \phi_c, \vec{p}, l_z} \sum_{p \in P} \left(l_p - l_z \frac{g(u_p, v_p, \theta_c, \phi_c, \vec{p})}{f(0, \theta_s, \vec{p})} \right)^2, \quad (2)$$

where P is the set of sky pixels and l_p is the observed relative intensity of pixel p in the image. Here, the nominator produces the zenith luminance value predicted by the sky model, and serves as a normalization term.

The atmospheric conditions are specified in the Perez sky model using five parameters \vec{p} . These parameters determine the appearance of the sky and can be related to turbidity t , which is a single parameter, commonly used to describe how clear or foggy the sky is. For low turbidities ($t \in [2, 6]$), i.e. between clear sky and thin fog [PSS99], the parameters \vec{p} can be mapped to turbidity t directly, allowing us to optimize over one parameter t , rather than five parameters \vec{p} .

To further reduce the number of unknowns, we take advantage of the increased availability of GPS systems and the EXIF data that is typically attached to the image. Using GPS data, time information from EXIF and a known solar model [RA03] we evaluate the sun zenith θ_s and azimuth ϕ_s angles. The image's EXIF data is used to compute the focal length in pixels, which is obtained from the focal length (usually given in *mm*) and the size of the camera sensor. Having the focal length allows us to make the optimization quite robust and makes the algorithm sufficiently stable to rely on the information obtained from a single image in comparison with images sequence [LGE10].

To minimize the risk of convergence to a local minimum we impose constraints on the sky appearance parameters \vec{p} or t , which is possible due to their limited range of values [PSM93, PSS99]. Following Lalonde et al. [LGE10] we initialize the camera zenith angle θ_c from the assumption

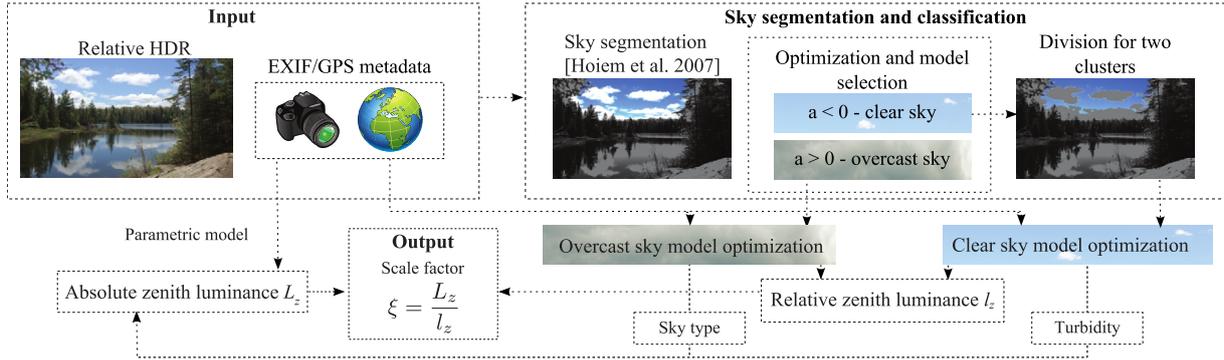


Figure 4: The flow of our algorithm. Note that when the input is an environment map the GPS/EXIF data are not required.

that the lowest sky pixel corresponds to the horizon. We also initialize the camera azimuth angle ϕ_c with a set of values in the range $[0, 2\pi]$ and choose the solution that gives the smallest residual. Note that especially when the camera is facing away from the sun, the accuracy of this solution is non-crucial (see Figure 1b).

Being a non-linear least-squares minimization problem with inequality constraints, (2) requires an iterative approach for its solution. We use the implementation of the trust-region-reflective algorithm from the MATLAB Optimization Toolbox as it fully satisfies our requirements.

2.2. Sky Segmentation and Classification

The first step of our algorithm is to extract portions of the image that depict sky. We segment the image using the segmentation method of Felzenszwalb et al. [FH04] and use the surface layout recovery method of Hoiem et al. [HEH07] to separate sky regions from the remainder of the image. Alternatively, the method of Tao et al. [TYS09] can be used to segment the sky region.

We wish to determine if this sky segment represents a clear or overcast sky. The segment's pixel data is conditioned by excluding 20% of the lightest pixels, as regions near the sun tend to be unreliable due to exposure problems. We note that one of parameters $a \in \vec{p}$ of the Perez sky model is indicative of whether sky pixels near the horizon are lighter than sky pixels away from the horizon ($a < 0$, indicating clear sky), or vice-versa ($a > 0$, indicating overcast sky). Therefore solving (2) for this parameter allows us to automatically determine whether the sky is clear or overcast and choose the optimal algorithm.

The necessity of having two algorithms is explained by the differences in absolute zenith luminance models, which we use for the scale factor estimation (1). For our datasets, we find that this approach correctly selects the type of sky for 95% of our clear images as well as for more than 80%

of our overcast images. In the case of clear sky images, the segmentation can be further fine-tuned using k-means clustering, whereby the cluster closer to blue is selected.

2.3. Relative Zenith Luminance

We are interested in recovering the zenith luminance l_z relative to the luminance values in the image. One can think about the value of l_z as the value that would be recorded if the image could be extended so that the zenith was visible.

Images with clear sky regions. The relative zenith luminance l_z for the case of clear skies is found as a solution of 2, where instead of optimization by \vec{p} the optimization is performed by turbidity t .

Images with overcast skies. Since the appearance of overcast skies is significantly different from clear skies, a different approach is required. The relative zenith luminance l_z estimation for overcast images is performed in two steps:

I. First, we solve a continuous optimization problem (2) to obtain an initial estimate of the model parameters.

II. Then, we estimate the type of overcast sky i_{opt} in the image following Kittler's classification for five types of overcast skies [KPD98]. The sky type is required for absolute zenith luminance estimation. Let \vec{p}_i be the set of sky parameters of one of the five overcast sky types, $i \in [1, 5]$. Then, the sky type i_{opt} can be taken as the one with the lowest l^2 -norm of the differences between its vector of sky parameters $\vec{p}_{i_{\text{opt}}}$ and the estimated optimal vector \vec{p}_{est} for a given image. Afterwards, the problem (2) can be solved again to find a better estimate of the relative zenith luminance l_z , with \vec{p} set to the values of sky parameters $\vec{p}_{i_{\text{opt}}}$ of the found sky type.

2.4. Absolute Zenith Luminance

To scale the image to absolute radiometric quantities, we also need an estimate of absolute luminance. For images with clear skies, the absolute zenith luminance value can be

obtained from Karayel's models L_z^{Karayel} [KNNS84] and one of Soler's models L_z^{Soler} [SG00].

We found that while both options give reasonable results, their accuracy depends on the sun elevation angle. To improve the accuracy over either model, we propose a weighted combination of L_z^{Karayel} and L_z^{Soler} , interpolating according to sun elevation angle $\zeta \in [0, 75]$ (degrees):

$$L_z = \cos^2(3\zeta)L_z^{\text{Karayel}} + \sin^2(3\zeta)L_z^{\text{Soler}}. \quad (3)$$

Based on the above interpolation, Soler's models [SG00] are prioritized for values between 15 and 45 degrees, whereas Karayel's model receives a higher weight for all other angles. Note that the interpolation weights always sum to 1. We validate this choice in Section 5.1.

For images with overcast skies the absolute zenith luminance L_z is computed from the CIE daylight distribution model [KPD98] using standard values of parameters for the appropriate overcast sky type i_{opt} .

3. Scale Factor Estimation for Angular Maps

In the introduction we give an example of how environment maps depicting the sky may be used to achieve realistic daylighting. In this section, we describe how our algorithm can be used to scale images of this form to absolute values.

The sky maps depict the whole sky hemisphere. Thus, the zenith sky element is visible in these images, potentially making the problem simpler. However, even if the sky has clear regions the sky zenith element could still be occluded by clouds. For overcast skies the visible sky element could not sufficiently represent the whole sky appearance. So, the optimization by the relative zenith luminance l_z could still be beneficial.

To perform the optimization we first need to express zenith angle θ and azimuth angle ϕ as functions of pixel coordinates. Here we give expressions for angular and latitude-longitude map formats. The conversions between different environment map formats can be found in [RHD*10].

Beginning with an angular map with pixel coordinates (u, v) normalised to the $[-1, 1]$ range, we derive polar coordinates $(\alpha, \beta) = (\pi\sqrt{u^2 + v^2}, \tan^{-1}(v/u))$. This allows us to express zenith and azimuth angles for each pixel as follows:

$$\theta_p = \cos^{-1}(\sin(\alpha)\sin(\beta)) \quad (4)$$

$$\phi_p = \tan^{-1}(\sin(\alpha)\cos(\beta)/\cos(\alpha)). \quad (5)$$

For latitude-longitude maps the zenith and azimuth angles can be evaluated directly from normalized pixel coordinates with $(\theta_p, \phi_p) = (\pi v/2, \pi u)$.

Further, we need to find the sun position. If the sun is not occluded its brightness typically exceeds the brightness of sky and clouds by a factor of fifty thousand. Correct maps should preserve this ratio. The technique of capturing bright sources is described in detailed by Stumpf et al. [STJ*04].

This ratio allows us to estimate the sun position as the point with the maximum intensity value in the image. Even if the sun is occluded by clouds this assumption will either hold or the environment will correspond to an overcast sky type where the sky appearance doesn't depend on the sun position. Then, sun zenith and sun azimuth angles could be found by substituting sun coordinates in the light probe into the above equations.

Estimating the sun position from the image itself allows us to lift the requirement of having GPS coordinates and time of image capture. Additionally, expressing the zenith and azimuth angles as functions of pixel coordinates only allows us to lift the requirement of EXIF data as well. Moreover, we do not require the camera direction. The visibility of the whole sky hemisphere allows us to obtain the angle γ_p between sun direction and sky element p 's direction from an image without knowledge of their absolute positions in the world. This enables the elimination of the camera zenith angle from the optimization. Thus, relative zenith luminance l_z can be estimated as the solution to the following problem:

$$\min_{t, l_z} \sum_{p \in P} \left(l_p - \frac{l_z}{f(0, \theta_s)} f(\theta_p, \gamma_p, t) \right)^2. \quad (6)$$

The absolute zenith luminance and the final scale factor can be found as described before.

4. Ground Truth Dataset

To validate our algorithm, we require a calibrated set of HDR images that are annotated with both EXIF and GPS data. To our knowledge, only the HDR Survey [Fai07] provides all this information. As this set contains a limited number of images with sky regions, we have created an additional calibrated HDR image dataset, annotated with EXIF as well as GPS location and orientation information. All images in this dataset contain a significant portion of clear or overcast sky. We plan to make this dataset publically available.

Our HDR dataset was photographed with a tripod-mounted Nikon D2h digital SLR camera, which acquires up to 9 images each spaced 1 EV apart, using autobracketing. The white point was set to a fixed 6700K (nearest to D65 that this camera supports) and the exposures were saved in the sRGB color space. Images were assembled into linear HDRs using Greg Ward's Photosphere application, which we have also used to derive the response curve of the Nikon D2h. We placed both an 18% grey card and a GretagMacbeth color checker in each scene for calibration. To scale the HDR images to absolute values, measurements of Yxy components of the grey card and several patches from the ColorChecker were taken with a Minolta CS100 color and luminance meter. We also transformed the images to the D65 white point, as per the sRGB standard. Using this set-up we collected 25 calibrated HDR images containing clear skies with varying turbidity and 9 images with overcast skies.

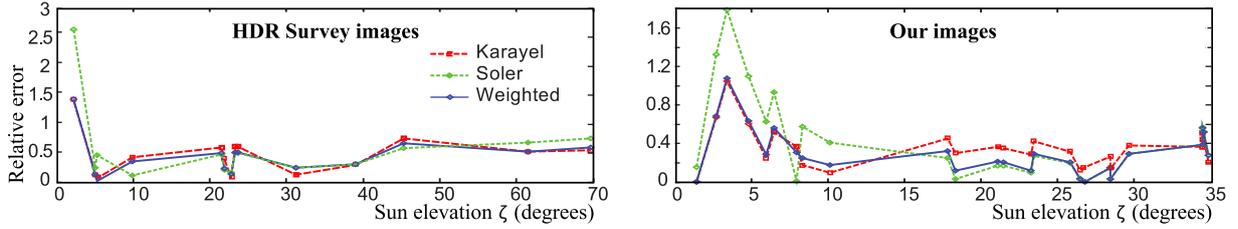


Figure 5: Relative errors for images with clear sky regions from the HDR Survey [Fai07] (left) and from our own dataset (right).

5. Algorithm Validation

In the following, we compare variants of our model for clear and overcast skies on a selection of images from the HDR Photographic Survey that contain sufficiently large areas of skies as well as on the images from our dataset. We use relative error in our comparisons: $\delta \triangleq |\xi - \xi_{\text{meas}}| / \xi_{\text{meas}}$, where ξ_{meas} is the scale factor taken from the HDR Photographic Survey or measured as described in Section 4.

5.1. Clear-Sky Images

The scale factor that allows the creation of absolute pixel values, is estimated using the various absolute zenith models (Section 2.4) to calculate L_z . The results are shown in Figure 5. We note that for sun elevation angles below 15 degrees Karayel’s model performs best, whereas Soler’s models perform better for sun elevations between 15 and 50 degrees. Our weighted combination of both models (Eq. (3)) accounts for this and produces errors smaller than can be achieved with either model. Using this weighted combination we obtain a mean error of 0.440 ± 0.344 (s.d.) for the HDR photographic survey [Fai07] and 0.309 ± 0.251 for our own dataset.

Note that though the measurements for indoor scenes can be taken with high precision, it is not an easy task for outdoor images. Thus, the measured values for some images in the HDR photographic survey could give scale factors different from each other by 40%. Additionally, for a sufficient number of images from Fairchild’s dataset the scale factor is estimated from measurements of a scene itself, which gives additional imprecision in ground-truth scale factor estimation. This can explain the higher mean relative error for images from the HDR photographic survey.

To demonstrate the performance of our approach we also compare our algorithm with techniques that solely rely on EXIF data. Such techniques compute the mean luminance of a scene as follows:

$$L = K \frac{N^2}{S t}, \quad (7)$$

where K is a constant, N is the f-number, S is sensitivity and t is the exposure time. In the general case these techniques



Figure 6: Two exposure sequences for HDR images of the same scene. The first scene is exposed for the background and the second one for the person standing behind. The exposure compensation value of the middle image in both sequences is zero.

Image	Mean luminance		
	EXIF	Our algorithm	Measured
background	12.5e+03	2.41e+03	2.95e+03
foreground	1.0e+03	2.59e+03	

Table 2: Mean luminance values for the images obtained from HDR sequences in the Figure 6, estimated using the EXIF-based approach (7) or by scaling mean luminance of an HDR image by a scale factor estimated using our approach. The estimates are compared with the mean luminance of the HDR images scaled by a factor estimated from measurements.

rely on several assumptions. The first is that the mean luminance in the area measured by the in-camera light meter corresponds to the luminance of an 18% card if this card would be inserted into the scene and measured. The second pertains to the used definition of sensitivity. The ISO standard 12232 [ISO06] includes several definitions of sensitivity: saturation-based ISO sensitivity and standard output sensitivity (see supp. materials). As discussed in [AG13] using the different definitions of sensitivity will give significantly different estimations for the same pixel. In their example they get two estimations, different by a factor of 1.43.



Figure 7: Example images from our dataset and HDR Photographic survey, which are used in the validation.

Further, the assumption about mean luminance can be easily broken when the spot/multi-zone metering modes are used. Let's consider typical candidate scenes for HDR outdoor images: portrait or trunks of trees on the background of sky. It may happen that the photographer will be interested in the foreground or background. This could lead to significantly different automatic exposure settings (Figure 6). Other sources of errors for exposure-based approaches are discussed in [WG07].

We compare our approach with the exposure based approach (7) as well as Ackermann et al. [AG13] modification of this equation. They suggest to account for the gamma curve of the sRGB color space. The radiance value L_p for each pixel p with intensity value I_p can be found from:

$$L_p = 15.4 \frac{f^{-1}(I_p/255) N^2}{f^{-1}(118/255) S_t}, \quad (8)$$

where f^{-1} corresponds to the inverse gamma correction of sRGB color space. This approach can not be used with HDR images, so we estimate luminance values for each pixel for the middle image of the HDR sequence, which in our case corresponds to an automatic exposure setting. From the absolute scaled values we then find the mean luminance of a scene. In order to get a scale factor for the HDR image, based on this approach, we divide the estimated absolute mean luminance by the mean luminance of the HDR scene.

We evaluate all the approaches on our clear sky dataset, where the image with the highest dynamic range has a range of 10 stops. All the images were captured using average metering mode. The images do not contain close up objects and are similar to the one in Figure 7c.

The value of K in (7) varies typically between 10.6 and 15.5. We keep K equal to 15.4 in (8) as suggested by Ackermann et al. We set K to 12.5 in (7) as the optimal value for our camera. We use the standard output sensitivity definition in both EXIF-data approaches as the one corresponding to our camera. Thus, our settings for equation (7) corresponds to the case of a fully known camera specification.

The images in our dataset meet the requirements imposed by EXIF-only methods. Under such conditions, EXIF methods yield correct absolute luminance values. The comparison of mean errors for our clear sky images is given in Figure 8, showing that our approach is statistically equivalent to the

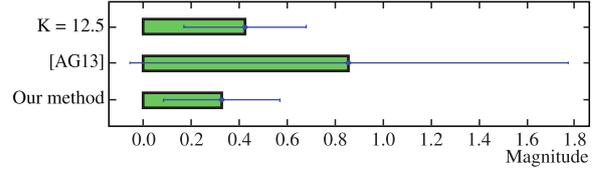


Figure 8: Mean and standard deviation of errors on images with clear sky regions from our dataset. Our approach is compared with [AG13] (AG13) and (7).

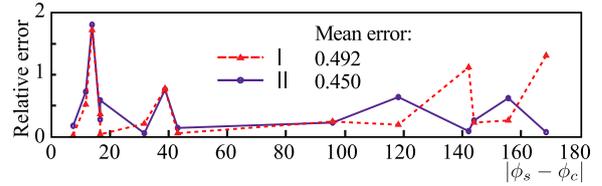


Figure 9: Relative errors for images with overcast skies as function of the azimuth difference between sun and camera. I corresponds to the case when we use $l_z \leftarrow l_{z_{est}}$ found in the first step. II corresponds to the full algorithm.

EXIF approach. This means that we deem our results to be accurate as well.

However, EXIF-only approaches are susceptible to error, for instance when other types of light metering are used. In the example of Figure 6 and Table 2 we show that the assumption about mean luminance can be easily broken for scenes with high dynamic range and depending on photographer intention. The same table shows that our method is robust against such variations, and therefore offers a significant advantage over EXIF-based methods.

5.2. Overcast Sky Images

Due to the limited number of overcast skies in either the HDR Photographic Survey and our dataset, we pool the results from both image collections. Relative errors calculated using step I of our approach as well as steps I and II (Section 2.3) are shown in Figure 9. This comparison shows that solving the optimization problem for a second time with sky parameters fixed to the parameters of the estimated sky type produces more accurate results, than using only step I.

5.3. Validation of Scale Factors for Angular Maps

To validate our algorithm on environment maps we used the sky probes from Stumpfel et al. [STJ*04]. The absolute luminance values corresponding to the intensity values in the sky probes are unknown, so the images were divided into two equal parts. The scale factor was estimated for a whole image, as well as its left and right parts. Then, the relative errors of the scale factors from left and right parts to the scale

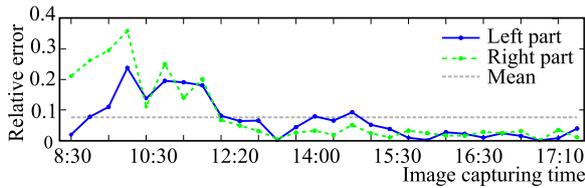


Figure 10: Magnitudes of relative errors of scale factors for the left and right parts of a sky probe to scale factors estimated from the whole sky probe.

factor from the whole image were evaluated. The mean error is less than ten percent for the full set of images (Figure 10). Results are in agreement with results of the validation of the algorithm on test images from our and Fairchild’s datasets.

5.4. The Hosek and Wilkie Sky Model

Recently, an extension for the Perez and Preetham sky models was suggested by Hosek and Wilkie [HW12]. This new model gives an improvement over the Perez sky model in modelling the sky appearance for lower solar elevation angles, as well as the appearance of the circumsolar ring. This is achieved by including two additional terms.

Moreover, the authors propose a new correspondence between turbidity values and parameters of the sky model, extending the range of its values from [2, 6] [PSS99] to [1, 10].

The introduced modifications require a zenith luminance model that is able to model the circumsolar ring. This is the reason why the authors, instead of normalizing by zenith luminance, find for each pair of turbidity and albedo values one “master value”, which depends on sun elevation angle. In this form, unfortunately, the model can not be applied to our algorithm. Instead, we require the normalization by zenith luminance as in (2), with the distribution of luminance values that two additional terms provide.

The results show that for our algorithm this model does not give advantages over using the Perez sky model (Figure 11). For images from our dataset, the Hosek model performs slightly worse than the Perez sky model. For images from Fairchild’s dataset, the Hosek model does not give any advantage in terms of mean of errors, but obtained results have a very low deviation from the mean. Given the increased number of required parameters, combined with the minimal improvements in our context, we have opted to retain the Perez sky model for our solution. The lack of improvement of the scale factor estimation using a new model is explained by the fact that we don’t rely on the regions where the improvement of sky representation was achieved. The region close to horizon is quite often occluded on the images. The region near the sun is automatically excluded on the sky segmentation step since without special equipment the SLR camera can not capture this region properly.

5.5. Discussion and Limitations

We validated our algorithm on images with diverse sky conditions and covering various sky areas (Figure 7). A limitation arises when the angle covered by the visible sky region is too small to contain sufficient information about sky intensity gradation. This is especially the case when the cloud appearance is highly complex.

It is often thought that veiling glare could have an adverse effect on radiometric measurements of HDR images. In our specific case, however, we analyse skies, which often exhibit a fairly uniform or low frequency luminance distribution. In addition, we typically measure high luminance values that are integrated over large image areas. As a result, we deem the influence of veiling glare on our measurements to be marginal.

Finally, the accuracy of scale factor estimation will increase if time-lapse sequences are used [LGE10]. Moreover, it could allow us to remove the requirement of known focal length and time of image capturing (information currently obtained from EXIF data).

6. Conclusions

We present an automatic algorithm to estimate absolute pixel values from single HDR images, using only the EXIF header information as well as GPS data. Note that, SLR cameras nowadays start to have built-in GPS systems and all recent mobile phones provide this tool. Conversely, the development of HDR technologies on mobile phones is a recent trend that makes our work increasingly relevant.

For images containing the full sky dome, as would be the case with light probes used for image-based lighting, our approach enables the estimation of absolute data without requiring any additional information.

Our method is targeted at outdoor scenes with HDR illumination, which are analysed to recover camera parameters as well as the turbidity of the atmosphere. Using a sky luminance model, the relative zenith luminance is estimated and subsequently correlated with the absolute zenith luminance calculated from GPS data. We compared several zenith luminance models and developed a trigonometric interpolation between two of them to obtain higher accuracy over a wider range of sun elevation angles. Additionally, we performed a comparison of the well-known Perez sky model with its recent extension [HW12] in the context of our work.

The algorithm was validated using an existing HDR image dataset as well as our own calibrated and annotated HDR image collection. We aim to make this dataset available.

We demonstrated that for scenes with high dynamic range simple EXIF-based approaches can be inaccurate. Limited to the images containing sky regions, our algorithm provides the solution which does not require any measurement equipment and can be used for certain existing imagery. Al-

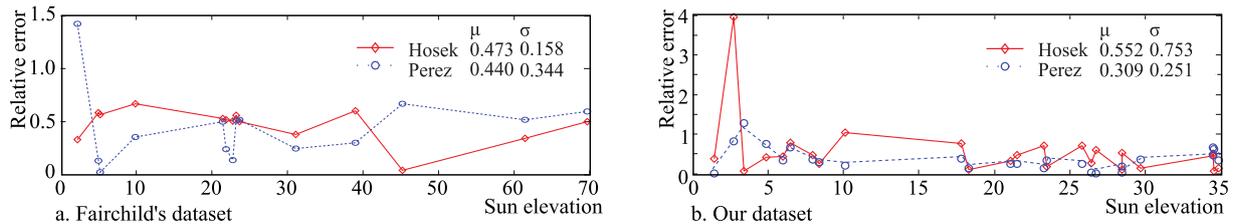


Figure 11: Comparison of magnitudes of relative errors for images with clear sky regions from Fairchild's dataset [Fai07] in cases of Perez [PSM93] and Hosek and Wilkie [HW12] sky models using weighted combination of Karayel and Soler zenith models. More comparison results can be found in the supplemental.

though, sky appearance is a highly complicated phenomenon and cluttered skies will be the difficult case for our algorithm, our algorithm is sufficiently robust as tested on the images with different sky types.

With the advent of scene-referred HDR displays, we think that there will be an increasing demand for images that are not only calibrated (i.e. linearized) but also absolute. Our algorithm helps create such data, and can therefore be seen as a tool that turns cameras into high resolution light meters.

References

- [AG13] ACKERMANN J., GOESELE M.: How bright is the moon? recovering and using absolute luminance values from internet images. In *Proc. of the 4th International Conference on Computational Color Imaging* (2013), Springer-Verlag, pp. 90–104. [2](#), [6](#), [7](#)
- [Deb02] DEBEVEC P.: Image-based lighting. *IEEE COMPUT GRAPH* 22, 2 (2002), 26–34. [2](#)
- [Fai05] FAIRCHILD M. D.: *Color appearance models*. Wiley, 2005. [2](#)
- [Fai07] FAIRCHILD M.: The HDR photographic survey. In *Proc. of the 15th IS&T Color Imaging Conference* (2007), pp. 233–238. [5](#), [6](#), [9](#)
- [FH04] FELZENSZWALB P. F., HUTTENLOCHER D. P.: Efficient graph-based image segmentation. *International Journal of Computer Vision* 59, 2 (2004), 167–181. [4](#)
- [FH09] FLIEGEL K., HAVLIN J.: Imaging photometer with a non-professional digital camera. *Proc. SPIE 7443* (2009), 74431Q–74431Q–8. [2](#)
- [HEH07] HOIEM D., EFROS A. A., HEBERT M.: Recovering surface layout from an image. *International Journal of Computer Vision* 75, 1 (2007), 151–172. [3](#), [4](#)
- [HW12] HOSEK L., WILKIE A.: An analytic model for full spectral sky-dome radiance. *ACM Trans. Graph.* 31, 4 (July 2012), 95:1–95:9. [2](#), [8](#), [9](#)
- [IM94] INEICHEN P., MOLINEAUX B.: Sky luminance data validation: Comparison of 7 models with 4 data banks. *Solar Energy* 52, 4 (1994), 337–346. [3](#)
- [ISO06] ISO: *Photography - digital still cameras - determination of exposure index, iso speed ratings, standard output sensitivity, and recommended exposure index*, 2006. [6](#)
- [KHFH11] KARSCH K., HEDAU V., FORSYTH D., HOIEM D.: Rendering synthetic objects into legacy photographs. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 157:1–157:12. [2](#)
- [KNN84] KARAYEL M., NAVVAB M., NEEMAN E., SELKOWITZ S. E.: Zenith luminance and sky luminance distributions for daylighting calculations. *Energy And Buildings* 6, 3 (1984), 283–291. [5](#)
- [KPD98] KITTLER R., PEREZ R., DARULA S.: *A set of standard skies characterizing daylight conditions for computer and energy conscious design*. US SK 92 052 Final Report, CA SAS Bratislava, 1998. [4](#), [5](#)
- [KZT13] KAWAKAMI R., ZHAO H., TAN R. T., IKEUCHI K.: Camera spectral sensitivity and white balance estimation from sky images. *International Journal of Computer Vision* 105, 3 (2013), 187–204. [3](#)
- [LGE10] LALONDE J.-F., G. N. S., EFROS A. A.: What do the sun and the sky tell us about the camera? *International Journal of Computer Vision* 88, 1 (2010), 24–51. [3](#), [8](#)
- [MN08] MADSEN C. B., NIELSEN M.: Towards probe-less augmented reality - a position paper. In *GRAPP* (2008), pp. 255–261. [2](#)
- [PSM93] PEREZ R., SEALS R., MICHALSKY J.: All-weather model for sky luminance distribution-preliminary configuration and validation. *Solar Energy* 50, 3 (1993), 235–245. [3](#), [9](#)
- [PSS99] PREETHAM A. J., SHIRLEY P., SMITS B.: A practical analytic model for daylight. In *Proc. of SIGGRAPH '99* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 91–100. [2](#), [3](#), [8](#)
- [RA03] REDA I., ANDREAS A.: *Solar position algorithm for solar radiation applications*. National Renewable Energy Laboratory, 2003. [3](#)
- [RHD*10] REINHARD E., HEIDRICH W., DEBEVEC P., PATANAIK S., WARD G., MYSZKOWSKI K.: *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*, 2nd ed. Morgan Kaufmann, San Francisco, 2010. [5](#)
- [RPK*12] REINHARD E., POULI T., KUNKEL T., LONG B., BALLESTAD A., DAMBERG G.: Calibrated image appearance reproduction. *ACM Trans. Graph.* 31, 6 (2012), 201. [1](#), [2](#)
- [SG00] SOLER A., GOPINATHAN K. K.: A study of zenith luminance on Madrid cloudless skies. *Solar Energy* 69, 5 (2000), 403–411. [5](#)
- [STJ*04] STUMPFEL J., TCHOU C., JONES A., HAWKINS T., WENGER A., DEBEVEC P.: Direct HDR capture of the sun and sky. In *Proc. of Afrigraph* (2004), ACM, pp. 145–149. [5](#), [7](#)
- [TYS09] TAO L., YUAN L., SUN J.: Skyfinder: attribute-based sky image search. In *ACM Transactions on Graphics (TOG)* (2009), vol. 28, ACM, p. 68. [4](#)
- [WG07] WÜLLER D., GABELE H.: The usage of digital cameras as luminance meters. *Proc. SPIE 6502* (2007), 65020U–65020U–11. [2](#), [7](#)