RGB Radiometry by digital cameras

Jan Hollan

Department of preventive medecine, Fac. of medicine, Masaryk uni., Brno

Abstract: *The author has developed a method of colour photometry using those common cameras, which are able to store raw data from their A/D converter. The resulting* raw2lum *programme is publicly available under the GNU license. The following text is an introduction to such photometry and a sketch of steps which are to be done to make a camera a measuring instrument, an imaging photometer.*

Contents

Introduction

Many disciplines need to "measure the light". Most often, it means "how much light comes to a unit surface", a quantity called illuminance, less often it means the luminance of various parts of an observed scene. The best known instruments for light measurement are luxmetres. Their function follows from their name, lux is the unit of illuminance. In spite of that, even luminance of some surfaces can be measured by them, approximately. Another classic instruments are so-called photomultipliers, very sensitive devices suitable for detecting faint amounts of light at night.

During the last decade of the second millenium, detectors based on CCD and CMOS chips became widespread .Their advantage is that instead of a single number they supply a matrix of numbers, linearly dependent on the amount of light hitting individual pixels during exposure. If the chip is located in the image plane of an optical system, then its output is a linear function of luminances or radiances (see more in Appendix 1) of the various elements of the scene, corresponding to the elements of a chip (pixels). In cases where an image capture brings an advantage, CCD or CMOS detectors are the most common scientific light detectors since many years.

In the third millenium, still-image cameras with such detectors became more common than cameras using 35 mm films. Digital cameras are no photometric instruments themselves. But considering the technology they are based on, they can be used in a way to become such instruments. The necessary condition for that is getting raw data from them, true values read out of individual pixels. Many modern still-image digital cameras fulfil this condition, being able to store date in proprietary formats with a generic designation RAW.

The author has offered (in 2002, with Petr Pravec, http://amper.ped.muni.cz/light/grant/) a project aimed at developing the needed software and beginning a photometric research of the night environment. The project has not been accepted, but the author began to solve it anyway. A breakthrough happened in 2003 thanks to a grant by the Czech Ministry of Environment (VaV/740/3/03, see the results at http://amper.ped.muni.cz/noc/), the needed software had been described and published within its final report in 2004 (http://amper.ped.muni.cz/noc/zprava_noc.pdf p. 25-27 and 37-39). Some years later, the software, esp. the basic programme raw2lum, is capable of using raw data (unlike jpeg or tiff files) from any camera and is available at http://amper.ped.muni.cz/light/luminance/.

The software development took place in the course of its application. Two studies for the largest Czech National Park (Krkonošský národní park) contributed to it in 2005 and 2006 (see the results in the http://amper.ped.muni.cz/noc/krnap) as well as measurements at another sites. Applicability and limitations of the method became evident. Ideas and applications of the method were presented at several conferences. The present paper summarises the method again, more compactly and up-to-date, moreover in English.

What do cameras register

Scientific digital cameras have broadband detectors, before which suitable spectral filters are placed according to the need. “Citizen cameras” differ by recording the image in three or four colours at once (RGB or YCGM in some old types). With one exception (Foveon chip, resolving colours by an increasing depth of capture with an increase of wavelength) the simultaneous three-colour image is achieved by having a colour matrix close to the chip, called Bayer matrix in most cases. Four pixels are needed to register “one colour point” (green pixels are doubled in the matrix). If the scene contains details with angular dimensions smaller than the 4-pixel element, the photometric information on such details is incomplete and any automated estimate of their R, G and B luminance may be misleading. Their photometry is made possible by defocusing (this is possible for bright stars on a much darker backround) or statistically by using many images which are shifted suitably (by one pixel roughly).

The colour matrix makes no obstacle for photometry of extended areas with uniform luminance. The only remaining problem is that standard programmes for imaging photometry assume that all pixels record the same colour. This can be overcome by converting the image to four “perforated” ones, each one in a unique colour (there are two green images, shifted by 1 pixel diagonally). I am aware of no software converting RAW to a series of four FITS (a format for scientific images) however.

I choose another way, namely developing an own programme which would use the RGB colours directly, as they can be stored in a pgm (portable graymap) format. The positions of the R, G and B pixels (or Y, C, G, M ones) within this file are given to the programme as a command-line parameter. A pgm file can be created from any RAW image by a famous programme dcraw (its early versions did not enable it, I had to patch the source code; I still use such a patch from 2003 for Fuji S5009 camera, which has the Bayer matrix rotated by 45°).

(A remark on non-RAW images: all cameras have a linear A/D converter producing numbers with 10, 12 or even 14 bits. But many refuse to store them, computing the assumed remaining two colours for each pixel by some interpolation. This operation changes the information, fabricates it. A further step is transforming the data by a secret non-linear way to an 8-bit representation during which the even the interpolated linear R, G, B values may become completely lost, if no reverse transformation exists or can be found. There have been authors attempting to use such images for photometry, hoping that the black-box transformations don't change or even calibrating each image separately – this is possible if there are defocused stars with known brightness in the image. Such processing is time-consuming however and the achievable accuracy is not good. I saw no reason to devote myself to such inferior hardware – actually, scanning images from film cameras might offer better, more reproducible results than a jpeg-based photometry. Cameras offering RAW files are widely available already, even if not within the cheapest category as a rule.)

It appeared that even RAW data from many cameras are subject to some unpleasant transformations. A common one is subtracting a constant which corresponds to a usual median or mean value of pixels which have not been exposed to light at all (in a so-called darkframe, an “image” taken with a covered lens). For a very underexposed image, what is a usual case for sky between the stars or for a night terrain, almost a half of the pixels are zero and carry no information about the amount of light incident at them. Another cameras, like Nikon D70, filter data by an unknown way, probably to reduce noise in the image. (Old Nikon cameras, which could be switched to a RAW mode by a special software, did not have this photometric fault). Having vast majority of pixels at non-zero values even in darkframes is a prerequisite for accurate faint light photometry. NEF format of Nikon and most Fuji cameras (apart from S2Pro) don't fulfil this demand. (Having to use Fuji S5000 anyway, I've incorporated a statistical compensation of this drawback directly into the raw2lum programme for this camera, see Appendix 2). Well-exposed image areas don't suffer from these photometry-unfriendly transformations, however, so all RAW-file cameras can produce good bright-light photometry when the eventual black-box transformations are minimised by choosing proper manual setting (or even by a special procedure during taking the images, to avoid an automated darkframe subtraction).

Luminance calibration in short

If a pixel value given by an A/D converter is linearly dependent on the luminance of the scene element projected to the pixel, as it is the case for CCD and CMOS instruments, it is easy to change the linear relationship to a direct proportionality, simply by subtracting a frame made with the same exposure time but with a covered lens (a darkframe). Even with no further calibration the relative change of luminance of the scene due to a change in its illuminance can be inferred taking another image. Either using the same exposure settings, or relying on the correctness of the camera-stored exposure information (it is stored in the so-called exif header of the file). The signal is proportional to the exposure, which can be computed as simply as “Exposure = ISO × exposure time / (1 s × f-number squared)”. My experience is that the accuracy and reproducibility of these values is very good, at the level of several per cent, if they are set manually (I found an exception for the shortest exposure times for Nikon 990, they had been longer in fact). In an automated exposure mode, some cameras use slightly another settings than they report in the image header, comparison of the images is not reliable in that case.

Any RAW images taken in the manual exposure mode can be calibrated later. The key step of the calibration is an image of a white surface whose luminance is known. The task can be divided into several phases.

The first one is finding such coefficients for R and B colours, that they would give the same value for white surface as the G colour (usually, this is the filter giving the largest raw values). These coefficients can be then applied to compute a proper colour image from the signal recorded through a colour matrix. We shall call these multiplied R and B values as *normalised*.

The second phase might be finding a coefficient by which the product of exposure and any of the normalised R, G, B values should be multiplied to get a proper luminance for white surface.

The third phase is finding a unite linear combination (sum of the coefficients being 1) which would yield a correct luminance even for another scene than a white one. An ideal linear combination would reproduce the spectral sensitivity of human vision. No real spectral sensitivities of camera R, G and B pixels offer such solution, the spectral sensitivity of photopic (daytime, foveal, cone) vision can be but approximated. A simple way of finding it is measuring luminance of surfaces of various colour by a luxmeter (or directly by a luminance meter), whose spectral sensitivity is similar to human vision, taking their images as well – the proper linear combination from the image values should give the same luminance. I've used another approach, namely finding the spectral sensitivities of individual camera colours and choosing such coefficients for these colours that the compound spectral sensitivity approximates the standard photopic one.

Luminance calibration in detail

Let me describe the real steps I applied for luminance calibration. Daylight is considered to be white, if the Sun is high in the clear sky. It is composed of stray light from the sky (in various hues of blue) and of direct sunlight (attenuated by dispersion in the air, predominantly in its shortwave, blue part, so having a yellow hue due to this attenuation). A colourless dispersing surface is needed as well. White paper is a possibility, unfortunately it contains an "optical brightener", a fluorescent agent converting the UV radiation to blue light. A surface painted by barium sulphate is a better alternative. A photometric white standard is the best diffuse surface – I had employed a Spectralon by Labsphere, which absorbs but one per cent of incident visible radiation. An image of a white standard under daylight is then a base for a normalisation of R and B to the G values. The raw2lum output gives medians of R, G and B in the chosen areas (tiles of the scene), the normalising coefficients are simply R/G and B/G for any area inside the white surface. (Quite possibly, tops of large cumulus clouds in summer might be an alternative to a white photometric standard...)

The second step, namely an absolute calibration of the camera, is evident, if the white surface is ideally dispersing and its illuminance is known. It *is* known if the Sun is hight in the clear sky. Sun is a very constant light source (its variations are seldom over 0.1 %). Cloudless atmosphere almost does not absorb light A small variable absorption of light is due to water vapour around 590 nm. Aerosol particles other than soot disperse light in the forward direction far more than backwards, so the sunlight dispersed by them hits the ground mostly. The variability of the ratio of direct and indirect component of sunlight (due to varying concentration of aerosol) plays therefore a minor role as far as a horizontal surface in an open landscape is concerned. Clear sky daytime horizontal illuminance is computed, e.g., by scripts generating the graphs at http://amper.ped.muni.cz/weather. Depths of shadows might be used to correct the basic estimate (even non-calibrated RAW images show them), or measurement of the so-called zenithal extinction. A series of images of a white surface oriented toward the Sun, spanning a large range of solar angular heights, can be used to compute that extinction (or a series of luxmeter readings, of course...). Knowing the zenithal extinction (attenuation of direct sunlight if the Sun would be in the zenith) enables to compute horizontal illumiance by *direct* sunlight with a relative accuracy of 1 %. Illuminance of the white surface can be directly measured by a luxmeter as well, common luxmeters have lower accuracy however. Of course they may provide a good check of the astronomically computed illuminance.

The luminance and illuminance of a surface are related by the so-called BRDF function, a ratio of luminance/illuminance dependent on the direction of incident light and direction of observation of the surface. Many surfaces can be considered as ideally dispersing if those directions are chosen suitably. It is rather coarse approximation for a white paper, a better one for a glazed white tile (neglecting the mirror-reflected component on its glasure), and a very good one for Spectralon. The luminance is then easily computed as a product of albedo and illuminance, divided by "pi steradians".

Even a less ideally dispersing surface can be used of course, if its luminance is measured by a luminance meter. An approximate luminance can be measured even by a luxmeter.

Another way of calibrating the camera absolutely needs a knowledge of the geometry of the image. Then the illuminance of the lens by a distant small light source can be computed (its luminance inferred from the image times its space angle) and compared with a measurement by a luxmeter. I did it imaging a window, with the rest of the scene having luminances hundred times lower. (As the light source, window was not white but

sky-blue, this was possible just after I got the spectral sensitivities of the camera colour pixels.)
Luminance calibration acquired by the above described procedures can reach an accuracy better than 5 %, down to a mere 1 %.

Spectral calibration

Taking many images in a laboratory with a costly equipment including a monochromator producing light with a known spectral radiance would be a way toward calibrating a camera spectrally.

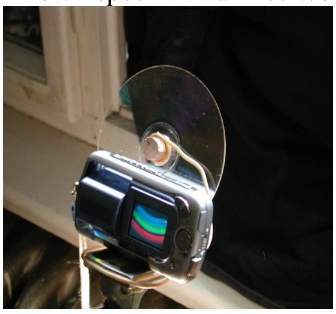

I used a completely different approach, imaging a solar spectrum produced by a DVD disk with no reflecting layer. Such a clear disk is sometimes available at the bottom of a stack of coloured or aluminised DVD-R disks. If the solar radiation hits any CD or DVD almost tangentially, the angular thickness of the Sun becomes irrelevant and the disk works as an excellent slit-less spectroscope. The curvature of the grooves in the disk makes an angularly short source appearing as a prolonged one, a curved line, when viewing the disk from a proper point (proper direction and distance from the disk). Dozens of dark Fraunhofer lines can be seen this way. The most prominent ones are apparent in the captured image. The registered pixel values can be then plotted as a function of their coordinate on the chip, selecting just pixels in some narrow band running perpendicularly to the spectral lines. Further, the usual formula for spectrographs is applied and its coefficients matching the coordinate on the chip to wavelength found by regression of known wavelengths of the observed solar spectral lines. The pixel values can be then shown as a function of wavelength: solar spectra as registered by R, G and B sensors (or YCGM in the example below, for Nikon 990).

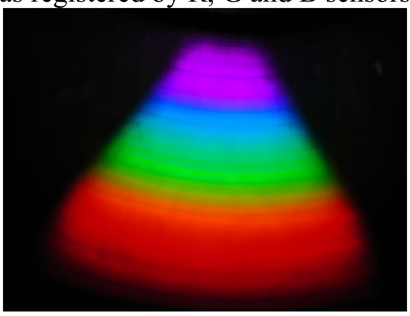
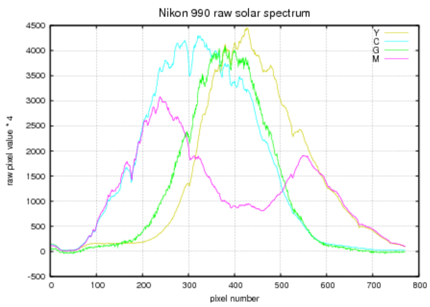

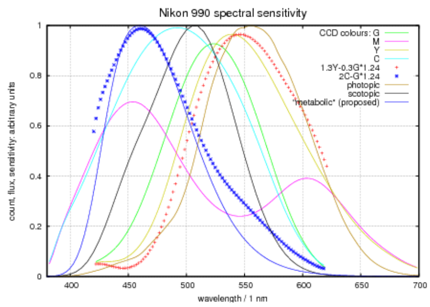

As said before, direct sunlight is not white, its intensity decreases toward the shortwave end of the visible spectrum. This decrease is known however, especially for the “standard atmosphere at airmass 1.5”. I took the spectral image at this airmass, I.e. at the moment when the Sun was near the angular height of 42 degrees in Brno, at 300 m above sea level. The raw solar spectra recorded by the camera can be then simply divided by the theoretic solar spectrum, and after some smoothing we get the wanted spectral sensitivities of each of the camera colours.

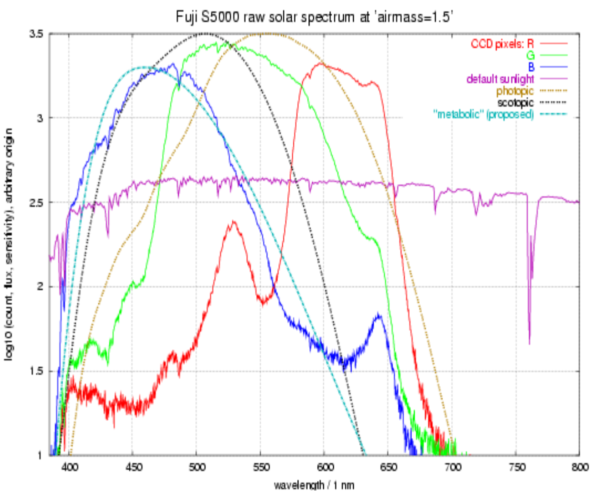

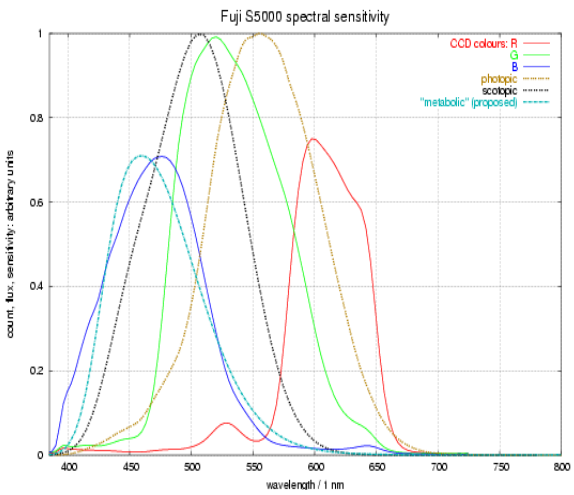

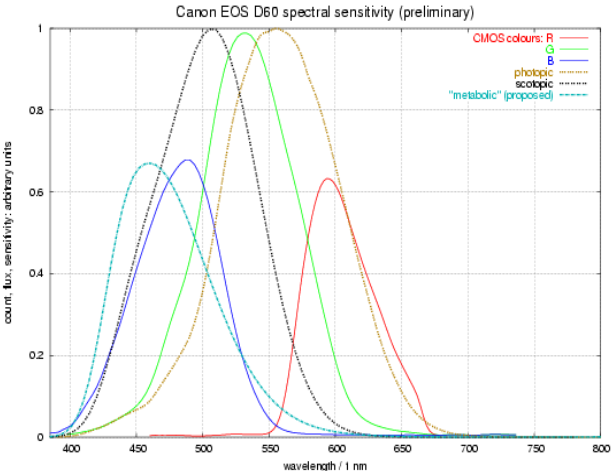

I have calibrated three cameras this way, Nikon 990 with a fish-eye converter, Fuji S5000 and Canon D60. Then I choose the unit combinations of normalised colours to match the photopic spectral sensitivity: Nikon 990: 1.3 Y − 0.3 G, Fuji S5000: 0.65 G + 0.35 R, Canon D60: 0.7 G + 0.3 R. The curves composed this way don't follow the photopic curve in details, but local relative differences are not over 30 % mostly (just for the Fuji S5000 they are). In an imaging photometry they could play some role only for light sources of unusual spectral composition or for conspicuously coloured surfaces. In any case, such a spectral-sensitivity based luminance calibration satisfies me more than guessing the coefficients just from a series of images of coloured surfaces whose luminance is measured by another instrument.

A second calibration task: flatfield or vignetting

The above procedures assumed that the dependence luminance-signal is independent on the pixel position within the image. But this is not the case, the signal is always lower in the corners of the image. This phenomenon is called vignetting or "light fall-off". A series of effects causes the phenomenon. If the angle of incidence is substantially different from zero, especially for wide-angle shots, the lens aperture diminishes for extreme angular distances from its optical axis (vignetting in the narrow sense of the word), transparency of the lens and Bayer filter matrix goes down, and even the absorptivity of the chip diminishes. In scientific photometry working with fields of view with small angular size the solution seems to be simple: taking an image of a scene with uniform luminance (so-called flatfield). Any other image is easily transformed to an image with a homogeneous luminance-signal ratio, dividing the difference of an object image from a darkframe by such a flatfield. It is not easy to find a scene with uniform luminance, but it is somehow possible for small angular fields of view. Flatfield-based normalisation solves also a problem of photometry of small details, uneven sensitivity of individual pixels and of shadows cast by dirt in the optical system.

Raw2lum cannot use flatfields. I did not find a way how to capture such an image with a large field of view, especially if its angular diameter is 180 degrees or even more (fisheye images). Taking several images through a pupil of a large integration sphere would be a solution, but I had no such a sphere. So I did not solve the problem of uneven sensitivity concerning details but just the problem of vignetting. I've employed images of a wall illuminated by a cloudy sky, changing the bearing of the camera. The goal was to find such a function of angular distance from the centre of projection, whose application would ensure that the computed luminances of the scene don't depend on where the camera points to. Actually, I solved the task just for two cameras in specific configurations: for Nikon 990 with a fish-eye converter and minimum zoom and for Fuji S5000 in two extreme focal lengths and one middle one, having minimum vignetting. In all cases, this was with full apertures (lowest f-numbers), diminishing the apertures would surely diminish vignetting as well.

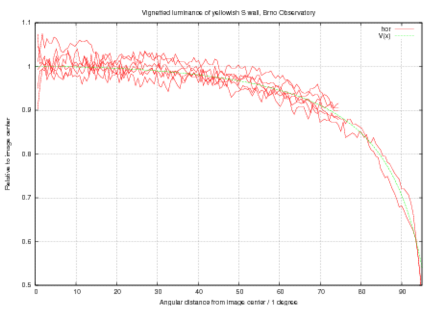

The third task: geometry of the projection

For images taken with a long focal length we may assume that the projection is the gnomonic one (central projection on a plane). This is the way how camera obscura works. Wide-angle lenses, and especially fish-eye ones cannot project this way, so the distance from the image centre is not simply $f \times \mathrm{tg}(z)$ where z is the angular distance of the point of the object from the optical axis of the camera and f the focal length. I've studied the geometry for Nikon 990 and the shortest focal length, using images of the starry sky. Adapting another programme of mine, map_bsct, I've created sky maps which matched the photographed sky. (Below, the negative image superimposed onto the map is shown at right, whereas in the middle, the colour-coded luminances are shown; see more on their coding below).

 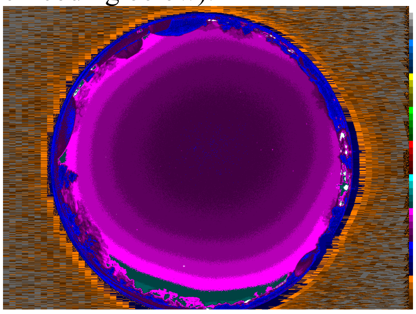 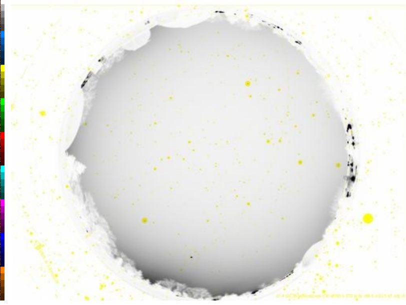

The additional projections I've arrived at are included in the published version of that mapping programme now. Radius from the centre of projection is proportional to the following functions of z:

equidistant from the pole of projection	z
conformal (stereographic)	$2\ \mathrm{tg}\,(z/2)$
gnomonic (central projection)	$\mathrm{tg}\,(z)$
equivalent (preserving space angles)	$2\ \sin\,(z/2)$
Nikon 990 with FC-E8 converter	$2.4\ \sin\,(z/2.4)$
Nikon 990	$1.4\ \mathrm{tg}\,(z/1.4)$
maybe for Pentax	$1.2\ \mathrm{tg}\,(z/1.2)$

Comparing the sky image with the map enables to find the focal length f, or generally the ratio of pixel modulus (well known mostly) and the focal length. In case of a fish-eye converter this was important, as the resulting values differed from the assumed ones.

Direct comparison with luxmeter readings

If the image geometry and vignetting is known and the corresponding parametres are incorporated into the software, another task may be solved using the luminance-calibrated camera. The illuminance of the lens plane by the imaged scene can be computed from the image, and the same value can be measured by a luxmeter independently. Using a full 180° fish-eye image, the two values should match, if the scene does not contain overexposed elements. If there are such elements, further images with 16× shorter exposure time each is to be taken. Then just the points with luminances within adequate limits are taken from each image and the illuminance of the lens plane by them is computed.

I made such a calibration check for Nikon 990 with FC-E8 converter, a system for which I've developed the software for images capturing a space angle over 2π srad. As said earlier, a similar check can be done under special circumstances, having bright object within the image and much darker scene behind the limits of the image, for any camera. If the scene surrounding the object is not dark enough, two luxmeter readings are to be taken, one with an unobscured sensor and the other one with an obstruction casting a shadow onto the sensor (so that just areas outside of the imaged part of the scene illuminate the sensor). The illuminance by the imaged part of the scene is the first reading minus the second one, and it should match the result computed from the image.

Imaging photometry by such cameras in practice

We have outlined all tasks which are necessary or recommendable to perform on any still-image camera which is meant to be used for photometry. If a RAW image is taken by such a calibrated camera in the manual mode (with all processing switched off if a camera offers any even for its “raw” images, which would not be raw anymore due to such processing) and a darkframe with the same settings follows or precedes it, all photometric data can be easily obtained by raw2lum from a pair of such files. In some cases, even the end of the text output of the programme may suffice, like that on an image of London sky (see ev. the overview of the 2004 photometry from London in http://amper.ped.muni.cz/light/luminance/london/)

```
File name     : ../28.jpeg
Date/Time     : 2004:09:09 23:53:23
Flash used    : No;    Focal length :  5.7mm
Exposure time: 2.000 s ;    Aperture     : f/2.8;    ISO equiv.    : 200
# Image percentils         L=(1.730*0.35*R+1.000*0.65*G+1.570*0.00*B)* 1.66E-4
# %     R     G1    G2    B readings  R         G         B         L / cd/m2
  1   -21      6     6   -38        -0.00603   9.97E-4  -0.00991  -0.00146
 10    28     56    56     8         0.00805   0.00930   0.00209   0.00886
 50    86    113   113    64         0.0247    0.0188    0.0167    0.0208
 90   140    167   168   115         0.0402    0.0278    0.0300    0.0322
 99   183    211   212   157         0.0526    0.0351    0.0409    0.0412
100   925   2751  5623  3538         0.266     0.695     0.923     0.545
# flux from 0.654 sr ( 8.59E-7 sr/4px) is  0.0131 lx ( 0.0148 lx tilewise);
# tilewise luminance  0.0244 cd/m2 with rel. SD 6.28 % (summing pixels below 15850)
# Fuji Devignetting: 1 # dark frame dark/32.pgm subtracted from all values
# Red  average=  85.0  minimum= -330  maximum=  925   over 15850:   0
# GrUp average= 111.9  minimum= -287  maximum= 2751   over 15850:   0
# GrDw average= 112.7  minimum= -267  maximum= 5623   over 15850:   0
# Blue average=  63.5  minimum= -324  maximum= 3538   over 15850:   0 zero:  1.3 %

End of results of "raw2lum" (version 1.22) with parameters:
  -c fuji -t 118 -d dark/32.pgm -fa -sm100:16 -i 28.txt -e 28.eps -n 28.ppm -s71:76
pgm/28.pgm
```

The important lines may be the just those on illuminance and luminance, saying that the sky luminance was about 24 mcd/m2 (as seen from the Thames near to London bridge toward city centre). Blue luminance computed just from the B pixels `(L=(1.730*0.00*R+1.000*0.00*G+1.570*1.00*B)* 1.66E-4)` is lower, as expected (London has orange lamps mostly); by “tilewise luminance” the mean value within rectangles sized some twenty thousand original R,G,B pixels is meant (areas large enough so that statistic compensation of zero pixel works reliably):

```
# flux from 0.654 sr ( 8.59E-7 sr/4px) is  0.0104 lx ( 0.0105 lx tilewise);
# tilewise luminance  0.0173 cd/m2 with rel. SD 6.73 % (summing pixels below 15850)
...End of results of "raw2lum" (version 1.22) with parameters:  -c fuji -t 118 -d
```

```
dark/32.pgm -fa -sm100:16 -i 28.txt -e 28b.eps -n 28b.ppm -nt -nh -MR1.730:1.0:1.570
-ML0.0:0.0:1.00 -s71:76 pgm/28.pgm
```

Visualizing the results is mostly preferable, especially for non-homogeneous scenes. False-colour coding of logarithm of luminance / (cd/m^2) may be the most useful way, a grid with data is another option, either as an overlay to this false-colour image or to the original image. The overlays contain coordinates of “tiles”, medians of G pixel values and, at their bottoms, the photometric results, luminances / (cd/m^2). The example shows the Big Ben and the sky above, the default logarithmic colour-coded luminance scale is given below the images (white codes the overexposed areas in case of the Big Ben clock, not the values over 200 kcd/m^2):

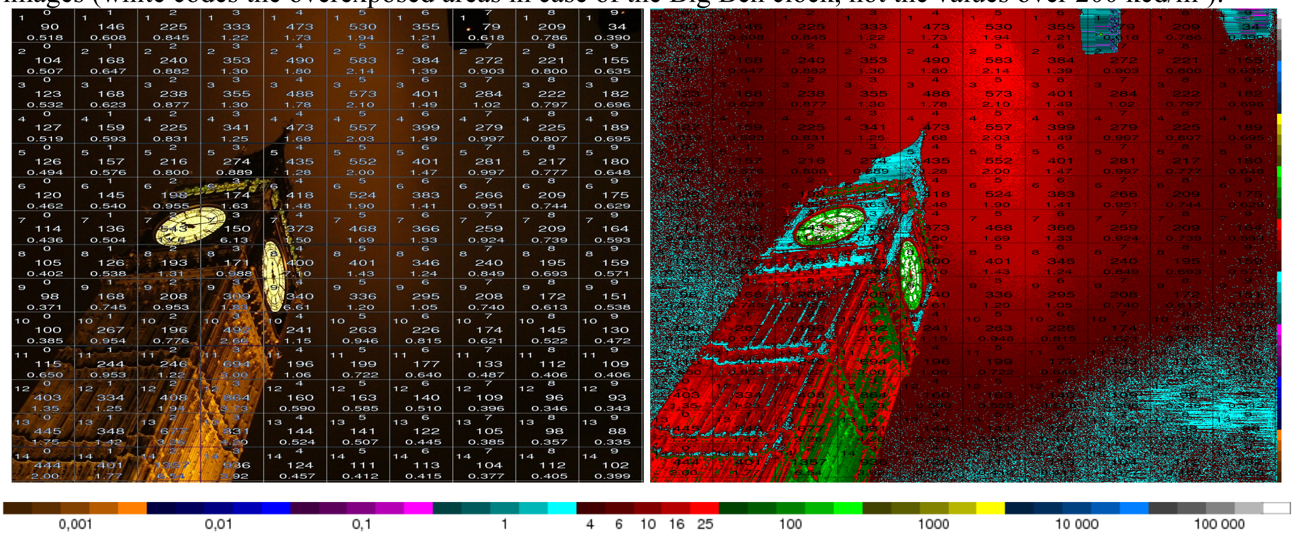

Cameras, especially those with a circular fish-eye view comprising 2 π srad at least (the full space angle is “four pi steradians” of course) can replace a luxmeter completely if needed, having an advantage of being able to measure much fainter light. And they are a much less expensive and much more versatile tool than professional luminance meters. Apart from that, as cameras perform multicolour radiometry, another photometric quantities than just photopic one can be obtained from them. Scotopic, of course, or “metabolic”, relevant for circadian rhythm of us and all animals. The latter case is shown above as a text output for the London sky (B luminance of Fuji S5000 approximates the “metabolic” luminance very well).

Many examples of photometric results obtained using still-frame cameras are available already, unfortunately not with an English description mostly. A rare case with a description is a photometry of a flame, http://amper.ped.muni.cz/light/luminance/candle/.

Appendix 1: Photometry and radiometry

Digital cameras, in a strict sense, offer just three-colour or four-colour radiometry. Calibrating them spectrally using a solar or another reference spectrum enables to report radiances in the colour radiometric system. These colour radiances are integrals of the object's spectral radiance over the spectral sensitivity of the camera, the result is to be expressed in watts per square metre and steradian. Any transformation to another radiometric system, like to photopic or scotopic luminances is but an approximation. It may be a very good one just in case that some of the camera's native spectral sensitivities matches the desired standard sensitivity curve rather well.

Photometric language, employing units descended from the basic SI unit candela and therefore from the sensitivity of human vision, is very convenient for colour radiometry at visible wavelengths. Within raw2lum, I've introduced a convention that for an insolated white surface the colour radiances are transformed to colour luminances in such a way that all of them are equal, the same as the photopic luminance.

In more detail: photometry defines but photopic and scotopic luminances, valid for daytime and nighttime vision, respectively. I've proposed (2004, Graz) a simple definition of another luminance, valid for receptors starting the day and night phases of metabolism and synchronising body circadian rhythm with the real day and night – these receptors form a non-imaging network of shortwave-sensitive cells. For white light, that “metabolic” luminance should be the same as the photopic one. (This differs from the scotopic luminance, which is anchored to match photopic one for monochromatic light at 555 nm and being several times larger for white light than the photopic one.) I propose the same procedure for R, G, B and further colour luminances, whenever they are inside the visible spectrum.

The only problem with the definition is, what should be regarded as a standard white light. An insolated horizontal white-painted surface gives even with the Sun hight in the clear sky somewhat variable spectral composition, depending on the air transparency. It is not much different from the CIE illuminant D65 however. Taken strictly, the luminances might be normalised to this CIE D65 standard.

Appendix 2: Compensating for the zero pixels

Many digital cameras "cut" the values obtained from the A/D converter at the level of a median of a darkframe. This poses a problem in photometry. Let us take an example we subtract a darkframe from a very little exposed image. Both frames differ just a bit, the difference being due to some photons captured by the exposed image. However, as almost half of the pixels remain at zero value, the captured photons can be "felt" in the non-zero pixels only, raising the sum of those many (millions) values somehow. Apparently, the real number of captured photons is proportional to the mean value of the difference of the two frames × (number of all pixels / number of non-zero pixels).

In an opposite case of a richly exposed scene it would be better to subtract no darkframe at all, as such a darkframe, if not rid of its bottom half (to be coded by negative numbers, if the median is set to zero), would have a zero mean value.

General images taken at night, having some parts underexposed, need some darkframe subtraction, however. The only solution I've found is subtracting the darkframe and adding some constant to each suitable set of pixels. For well exposed ones, it should be the mean value of the darkframe. Less exposed parts, being below percentil 99 of the darkframe, should be added "something less". I'm adding an integral of the function "percentile value" up to the percentile whose value matches the mean value of the given area. If the mean value of the area is larger than that corresponding to the 99. percentile of the darkframe, the integration ends at this percentile – the result being almost the same as a mean value of the darkframe.

This procedure is applied within raw2lum just for a single camera, Fuji S5000, for which I have approximated the values of finite integrals from zero to the given percentile by a polynome. Of course, such a compensation for zero pixels can be applied in a statistic way only, it works entirely well for regions with a uniform luminance consisting of thousands of pixels. I'm employing it, however, even for groups of 16(×4) pixels in a colour-coded logarithmic representation of luminance (each 4 original pixels forming just one element of the luminance-showing image), as it improves the false-colour coding of luminance a lot. More on the problem and its solution within http://amper.ped.muni.cz/noc/krnap/2006/low_exp.htm.

References

Hollan J: RGB radiometrie digitálními fotoaparáty. Brno 2006, ... Available at http://amper.ped.muni.cz/light/luminance/czech

...to be continued...