RAW$_d$

An extension on RAW mode images

Michiel Kamermans
www.nihongoresources.com

July 30, 2008

Abstract

Traditional image formats are static, one-time snapshots of a scene, specifying color luminescence in some color space (RGB, HSV, etc). However, storing photographic images in this fashion misses out on some of the lighting information that can be derived given the technology used. For the rest of the duration of this brief I will discuss an alternative image format, applied to pixels in a digital camera's optical recording chip (not to be confused with computer screen pixels).

1 introduction

The use of traditional photography using analog media has led to an entrenched belief that photographic images capture a moment in time. However, this has only been the case because of the medium used: the analog film roll works because it works as an accumulator, not so much recording how much light came in on a particular point of the photosensitive surface, but recording the total accumulated amount of light per reactive point by virtue of photochemical reaction. While this is the nature of analog photography, there is a genuine step forward that has been missed with the move to digital photography that I would like to highlight in this writing.

Digital photography, uniformly, has tried to recreate the analog experience by offering all the familiar controls of analog cameras, with the only real addition the ability to "see what you did" as you do it. While we could consider in-camera adjustments like algorithmic sharpening, noise reduction and white balancing great improvements, they are nothing that could not already be done for analog photographs using any number of widely available software packages. Moving this functionality into digital cameras has not so much improved photography, as it has reduced the need for several remedial tasks that come with being active in the field of photography, be it as amateur or professional.

However, with the move from analog to digital photosensitive media, one powerful feature of the medium has been overlooked: the new medium technically allows us to record a scene in a way that is genuinely different from the accumulate approach. Rather than recording the accumulated luminescence of some point at time interval X, we can actually record the accumulation gradient itself. Rather than recording the result of light hitting a point on the photosensitive surface for a time interval X, we can explicitly record the function that tells us how much light per time unit hits the points on our digital photosensitive surface.

Using digital photography, we can record how a scene evolves during the shot interval, rather than merely handing us the end result of that interval. It is difficult to overemphasize how big a step forward this would be for photography.

2 The theory

My suggestion is to move away from the RAW mode images as they are shot now, where for each camera pixel an accumulated luminescence is recorded over an interval X, instead moving to an image format that specifies the linear function y=a+bx that describes the luminescence gradient over time, given some initial luminescence offset. If a photosensitive chip (be it CCD or CMOS) is sampled several times during an exposure interval, then it is possible to - with a minimum of processing - abstract the luminescence function for each camera pixel in such a way that, if evaluated for the recorded exposure interval, the same image would be obtained as when the shot was made in conventional RAW mode.

While using the function to render the pixel's data means that the resulting image will look the same as a RAW image when it is rendered with the same exposure time as the original shutter speed dictated, it offers a world of new functionality in post production: with access to the luminescence function, an image (or even just selective parts of an image) can be 'developed' with alternative exposure times, allowing fine grained dynamic lighting control - something which (perhaps bizarly) currently has to be done through interpolation of several images of the same scene shot at different exposure times to prevent over-exposure on highlight regions and under-exposure on shadow regions.

With access to the luminescence gradient, these two problems would simply cease to exist. With access to the luminescence gradient rather than the absolute accumulated luminescence, it would be a trivial matter to pick exposure times that compensate for the fact that some (or all) part(s) of an image would be over- or under-exposed.

3 The procedure

The procedure to obtain a RAW_d image is essentially very simple. The camera is told to take a picture, and as light starts being recorded by the photosensitive chip, an initial 'base value' is determined for each pixel (giving us the offset for our linear function). While the shutter has not yet closed subsequent samplings for each pixel are performed so that an intensity delta can be computed (giving us the luminescence coefficient over time), and a pixel in RAW_d mode is described by the function a+bx, where 'a' is the initial luminescence offset, 'b' is the luminescence coefficient, and 'x' is the time component specified in some unit of time. Reconstruction, or rendering, of the image would then consist of evaluating the luminescence function for the appropriate time interval length 'x' for each pixel in the image.

4 Potential pitfalls

It is of course necessary to pick a time unit that ensures a luminescence change of 1 can be recorded, which highlights one challenge for implementing this image format: if the fastest shutter speed is 1/4000s, for instance, then this corresponds to 250 nanoseconds of exposure. If we sample the chip N times, then under these condition, any pixel that does not increase luminescence by 1 over at least 250/N nanoseconds will have a - quite likely incorrect - luminescence delta of 0, rather than some small value between 0 and 1. Thus, it will be necessary to pick a 'useful' time unit and record this along with the image data. This is easily doable per individual camera, where the recording properties of the photosensitive chip and minimum shutter speed are known, but may also be possible on a per-shot-image basis, although in such cases it would lead to a certain amount of processing overhead that might slow down the FPS rate of a camera employing this RAW_d imaging mode.

In addition to this, responsibility for indicating when using RAW_d format would be unsuitable, and switch over to normal RAW mode instead, would lie with the camera itself. If a shot is set up in such a way that the photosensitive chip will be unable to perform the desired number of samplings with which to determine the luminescence function, then the initial offset recording step can be skipped and the final sampling of the pixels on the chip will give the same result as a conventional RAW image would. However, as this would only be the case at extremely high shutter speeds, and the fact that – with most every major player in the digital photography world moving to CMOS sensors (which have very fast readout indeed, as is evident by CMOS chip use in 1000+ FPS cameras) – this problem will likely manifest in only a small number of photographic circumstances.

5 The pro's and cons

Advantages

- Equivalent to RAW format when rendered with same exposure.
- Backward compatible with RAW mode images, by setting real data as offset with a gradient of 0.

- Virtually impossible to under- or over-expose an image.
- Allows far greater control over image lighting than conventional RAW mode.

Disadvantages

- Likely a greater file size than conventional RAW mode images.
- Requires determining a suitable time unit, either per camera or per image shot.
- Requires additional circuitry on chip, or additional processing by the camera OS, to compute the luminescence gradient per pixel, affecting the number of frames per second a camera can shoot in this mode.
- Number of samplings of the photosensitive chip may be restricted at high shutter speeds.

6 Analysis

Analysis of the properties of RAW_d mode images:

- Underexposure due to too small a sampling step.
- Overexposure due to pixewl blowout in a single sampling step.
- Larger than RAW images, due to the storage of the gradient value, expressed as integer value increase per single exposure step.
- Loss due to rounding when considering gradient per single exposure step (i.e. 7:2 would be 3.5:1, but would be integer stored as 3:1).
- Reduced FPS due to processing overhead.

7 Conclusions

Conlusions go here. Basically I would advocate taking RAW_d mode images over conventional RAW mode images any day: the size difference is a non-issue if you own a "phototank" or have some way to dump your image data to from time to time. Given the enhanced image properties you get in return, it's just not worth arguing file size.