The Future of Digital Imaging - High Dynamic Range Photography

Jon Meyer, Feb 2004

[Updated April 7 2005 in response to this thread]

A new crop of technologies using "High Dynamic Range" will push digital imaging way beyond traditional media...

> Newsflash: Adobe has announced Photoshop CS will support High Dynamic Range imaging in their next version, available in May 2005. (details)

Today's digital cameras match or slightly exceed the performance of silver halide film. Computer graphics has achieved the goal of photorealism. Now is the time to go beyond simply matching paper and silver halide, towards perceptual displays that can show anything that our eye is capable of seeing.

One area where current digital imaging technology lags is *dynamic range*. In this article, I'll look at recent advances in the field of high dynamic range imaging (HDR). I will cover the basic concepts of dynamic range, and talk about new HDR technologies.

Caveat: I am a photographer, approaching HDR imaging from the perspective of its expressive potential. I am not involved in HDR research, marketing, product development, etc. Neither am I a perception or imaging expert. The information given here is almost certainly incomplete, and may contain innacuracies.

The Problem

Here's the problem in a nutshell:

- Real-world scenes contain light ranges that exceed a 50,000:1 dynamic range.
- For over a thousand years, media has been limited to around a 300:1 dynamic range.
- So you have a mapping issue: how do you represent light values in a scene using a much more limited set of light values for a particular media?

If you are not careful, you end up with results like this:

You can see the chair but nothing out of the window.

Old Solutions

Master painters were very clever about mapping scene intensities to canvas. They used a large number of tricks. Look at El Greco's *La Agoria en el Jardir* from 1590:

El Greco used saturated colors of opposing hues to increase the apparent dynamic range of the scene. He also painted black or white lines around the edges of contours. Our eyes determine contrast locally, so increasing the contrast at local edges increases the overall perceived contrast of the scene.

(Generally, "dynamic range" refers to the measured ratio between high and low extremes in a set of values. The word "contrast" is often used interchangeably with dynamic range. I prefer to use "contrast" to refer to the perceived contrast of a scene, which may be different from the measured dynamic range. El Greco's edge contours increase perceived contrast but don't change the dynamic range of canvas).

Here is another example from Monet. The red sun in his famous *Impressions at Sunrise* really leaps off the canvas (maybe not on your web browser!). The sun is actually the same brightness value as the surrounding clouds. However, since the sun is a saturated red placed over a saturated blue cloud, it creates a color vibration, and the sun looks much brighter than it is:

Check out webexhibits way-cool Flash illustration of this.

Dynamic Range in Photography

Photography is the act of capturing a scene with an imaging device and outputting the resulting image on a display device.

The dynamic range of the imaging and display device play a huge role in achieving good results. If the dynamic range of the source scene is too great, sacrifices must be made: either the shadows go, or the highlights go. Photographers have to know and work withing these limitations.

W. Eugene Smith spent five days in the darkroom until he came up with a print of Albert Schweitzer that he was happy with:

Smith was dealing with the issue that silver halide negatives have a greater dynamic range than photographic paper - so he had to "dodge and burn" different areas of the image to get a result where both the lamp and sitter are visible.

Perhaps the greatest master of dynamic range in photography was Ansel Adams. He was the first to systematically measure the sensitivity range of all of the equipment he used. His "zone system" let him predict precisely what details he could capture on film and paper, so he could make decisions before pressing the shutter:

Color Photography

Color negative films have less dynamic range (or "latitude") than black and white films. All the layers and dies reduce the sensitivity of the film. The first color films had very poor latitude, so film manufacturers added more layers - each color layer was split in two, a high-sensitivity and a low-sensitivity layer, using different crystal formations:

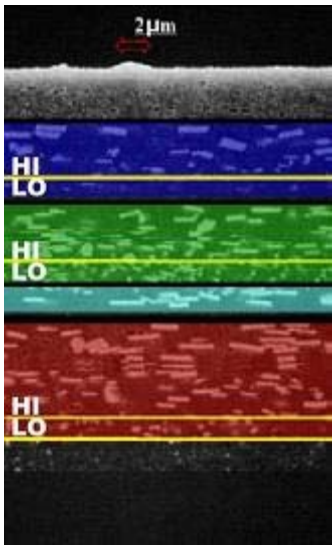

(I'm not an expert, but maybe color positive film doesn't use this trick, hence the difference in latitude between positive and negative film?)

One way to get really great dynamic range with color photography is to use black-and-white film and color filters. You have to take three photos, one with a red filter, one with a green, and one with a blue - and then compose the three images together. If you use glass plate negatives, you end up with images that have incredible colors and resolution. See below:

The most amazing thing about this photo is that it was taken in around 1915 by Prokudin Gorskii. And while it is true that this image was digitally

enhanced in Photoshop, my own experiments with a 4x5 and three filters alsc produced a distinctly non-kodachrome palette.

The Next Horizon: Digital HDR

Over the next decade, the imaging industry will inevitably transition to high dynamic range (HDR) imaging, creating devices that provide a latitude range far greater than traditional silver halide film. This change will affect all aspects of image making. Each of the systems in the image workflow will be modified, including capture, storage, editing and output.

Let's look at each of these workflow stages in turn.

Capture

Today's cameras have ample resolution. The next area product differentiatior that camera manufacturers will compete on will be the quality of the pixels, not simply the number of them. This is starting to happen already.

Consumer cameras may lead this drive, and be instrumental in driving forward the change to HDR. It is interesting that Foveon is currently courting Polaroid, for example.

Fuji's SuperCCD S3 Pro camera has a chip with high and low sensitivity sensors per pixel location to increase dynamic range (sound familiar?).

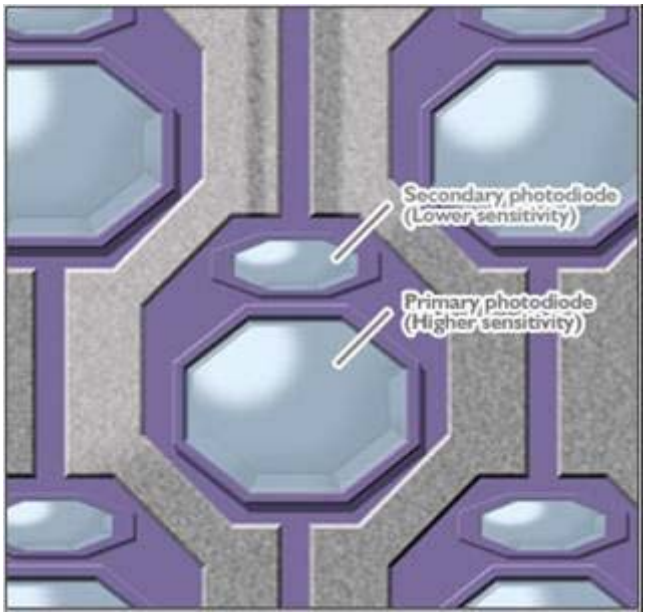

Although the resulting chip has lower overall resolution, it captures greater dynamic range. This tradeoff of resolution for dynamic range is the beginning of an important trend.

An alternative approach is to merge multiple images to increase dynamic range. Paul Debevec at SIGGRAPH 97 showed how to take multiple photographs at different exposures and merge them together to create a single high-dynamic range image. This technique is now incorporated in products such as Photogenics. For now, the technique only works well if the camera is mounted on a tripod. Researchers have already built HDR image "stitchers" which merge multiple images and automatically account for camera motion between snaps.

For most consumers, "HDR" will simply mean that the camera records more details in shadows and in highlights. Just as RAW images pushed the envelope of what you can extract from a digital image, HDR will further increase the available tonal range.

Consumers will benefit from the true point-and-shoot ability that broader latitude offers, because HDR cameras will produce usable images from a much wider range of lighting situations.

Eventually point-and-shoot cameras will lose their built-in flash. Anyone who uses a camera with a cheap flash soon learns that the pictures generally look better if you turn the flash off. Sensors are becoming more sensitive, cameras

are getting smaller, and light metering is getting smarter. Add three or four stops of dynamic range, and the flash becomes a creative ad-on rather than a requirement.

Professional photographers will also benefit from HDR. With HDR technologies, photographers can really push the creative envelope, exploring the extremes of high-key and low-key effects.

Professional cameras will offer a multitude of HDR image-taking modes. For example, they will automatically blend multiple images taken with different exposures, with and without flash, possibly using multiple light sources, to produce a single and extremely maliable master image.

Storage

All image file formats have range limitations. Formats such as JPEG and GIF provide eight bits per color channel (often referred to as 24 bit images). These are ideal for low dynamic range images. You can represent a linear range of intensities from about 250:1. In practice, JPEG uses a gamma curve rather than a linear range. Although this extends the dynamic range of the format, it does so by reducing the smoothness of the steps. JPEG is great for display, but once you start editing the image, you start seeing artifacts caused by the limited range of values.

Newer formats including JPEG2000, RAW and PNG offer up to 16 bits per color channel, which is plenty for most purposes. However, there is no support for "underage" or "overage" - these image formats state that "0" should be mapped to the darkest black of the display, and "65536" (or the equivalent) should be mapped to the whitest white. If you want to represent images that contain brightnesses beyond what your monitor can currently display (e.g. as produced by HDRShop), you need to look elsewhere.

The most exciting HDR image format today is OpenEXR, developed by Industrial Light & Magic.

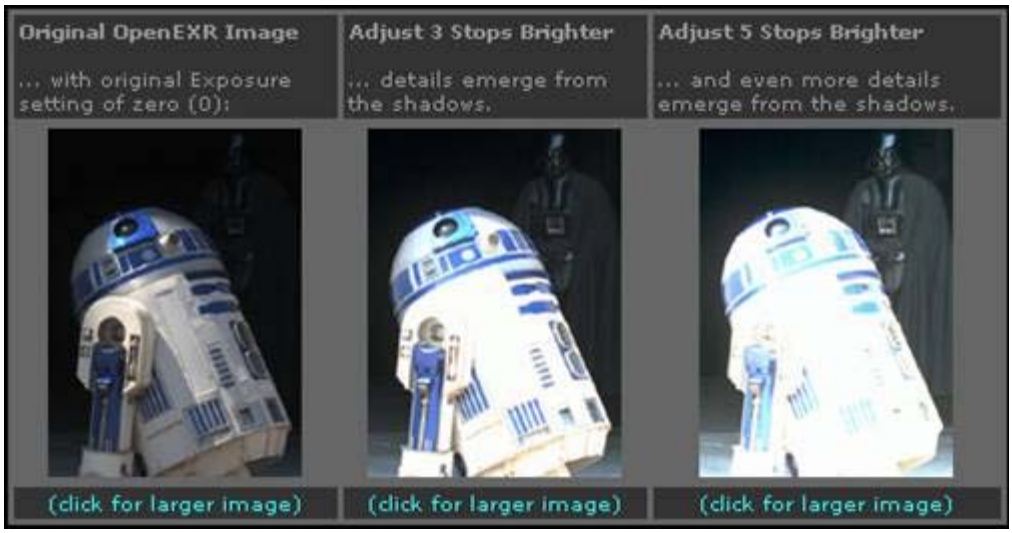

I say this partly because their documentation includes photos from Star Wars (see above). But it also supports both 16-bit and 32-bit float representations, lossy and lossless codecs, and has a great definition for underage and overage

Other examples of High Dynamic Range formats include SGI's TIFF LogLuv format, Floating Point TIFF format, Radiance's RGBE format, and Portable Float Maps (PFMs).

HDR image formats are especially significant for archival and stock uses, since they store data with enough precision to record what we can see, rather than what our displays can show.

There are a range of proprietary formats that offer medium or high dynamic range. The various RAW formats support whatever dynamic range the underlying device associated with the RAW file uses. Personally, I am not a fan of RAW formats for long term image stoage, since they are too device and vendor specific. However, that's a whole other debate.

I don't yet know if OpenEXR will become a consumer standard, or if it will remain a file format used only by Hollywood. Microsoft, Apple, Adobe, Canon and others will no doubt have a big hand in shaping that decision.

Editing

Of all imaging tasks, editing is the one that demands the highest dynamic range. Editing operations need high precision to avoid aliasing artifacts such as banding and jaggies.

Audio professionals know this. Editing tools like ProTools already use 48 bits per sample internally, even though the common CD output format only supports 16. Why should we image makers accept anything less?

Recently, Photoshop announced supports for 32-bit images. I've known this was coming for a while, and am very excited this has finally been announced. But I don't have a copy of the latest Photoshop (May...), so I don't know if the 32-bit support is as limited as the 16-bit support is today, or if they have done more work to add 32 bit support to all the filters and tools. Either way, its a fantastic step.

Idruna

Idruna Software is another company doing interesting HDR software. I played with their PhotogenicsHDR when it first came out, but I found it a little hard to use. Perhaps the newest PocketPC version is different...

Photoshop users are all familiar with the issues of low dynamic range. With 8 bit channels, if you brighten an image, information is lost irretrievable: darkening the image after brightening it does not restore the original appearance. Instead, all of the highlights appear flat and washed out. To avoid this problem, you must work in carefully planned stages, using complex arrangements of adjustment layers.

With a true HDR tool, if you brighten an image and then darken it, you would see something very close to the original image. True HDR editing tools will

enable image workers to follow a much more flexible and simplified work approach, using fewer adjustment layers, and with much less worry about aliasing artifacts. I expect HDR software will lead to increases in productivity and greater expressiveness.

It will take the imaging software industry some time to retool and retrain. With HDR, for example, you run into the issue of representing brightness values that are present in the image but beyond what your display can show you. Also, creating a color mixer that works in floats is hard, since there is no longer a "white" - instead you have a "paper white", "lightbulb white", "duluxe brilliant white", "sunlight white", and so on.

Another unsolved issue is size: If each channel of an image is 32 instead of 8 bits, the image becomes four times larger. Switching to HDR therefore makes a 100mb image take up 400mb. Not surprisingly, editing operations take about four times longer. Software will need to become smarter about scheduling work. Live Picture, an early image compositing tool, did a good job of this, but is no longer available. I expect to see a revival of these techniques as people grapple with 10GB images.

Tone Mapping

If you want to print an HDR image on paper or show it on a LCD panel, you must somehow convert the wide intensity range in the image to the lower range supported by the display. This process is called tone mapping.

One solution is to use manual tone mapping - i.e. the familiar dodge-and-burn techniques used in silver halide photography. HDR software will of course support manual dodge and burn.

Another solution is to use a tone mapping filter. There are already several to choose from.

The right image above show's *Fattal, Lischinkski, and Werman's* tone mapping algorithm filter applied to Debevec's HDR photo of Stanford Memorial Church. The left image shows what you get if you display the image without tone mapping.

This is a hot area of research in computer graphics. As with HDR file formats, there is currently no clear winner. Several techniques are listed in the resource section. I expect the major companies to each champion their own tone-mapping technologies in service bureaus and print finishing.

In addition, there are already a number of smart tone mapping filters, which automatically reduce the dynamic range of an image.

Output

There are already several displays on the market offering extended output range. The most astonishing is the Sunnybrook HDR display with a claimed contrast ratio of 60,000:1, good enough to reproduce the effect of a sunlit scene. They achieve this using high-power white LEDs.

The only bad thing about the Sunnybrook display is that once you look at it for a few minutes you just assume that this is how images are supposed to look - it is such a transparently great technology that until you see a normal image on a normal display you don't really think of the HDR display as that exciting. The display is still in the very-expensive bracket, but this will change quickly.

Of course, to fully make use of the Sunnybrook display, you need lots of HDR images. I anticipate a huge market for libraries of stock images in HDR.

Applications

Today, the main users of HDR imaging devices are specialized professionals working in the film, animation and VR industries. Some applications are listed below (see the Resources page for links).

Films - Tools such as HDRShop by Paul Debevec enable you to convert a series of photographs into a light probe - a special image that represents the lighting environment in a room. You can then use the light probe to light virtual objects, so that the virtual objects actually appear to be lit by the light from the room. This technique is especially useful for compositing computer graphic objects into images of real scenes. Hollywood films use light maps extensively to blend CGI into a scene.

Panoramas - Another use for HDR is in panoramic images. Panoramas often have a wide dynamic range, e.g. one part of the panorama may contain the sun, another part may be in deep shadow. Online web panoramas constructed from HDR images look much better than non-HDR equivalents.

Computer Games - A third use for HDR is in computer games. Recent computer graphics cards support HDR texture maps. With HDR texture maps, you can render objects using light probes, in real time, yielding much more dynamic and interesting lighting effects.

As more consumer-oriented HDR products arrive, I expect that the largest application of HDR will be in consumer photography.

Do we really need HDR?

I recently read this comment from Sam Berry:

> ... the whole article has no mention of the fact that the reason most controlled lighting is almost always done to ratio of less than 8:1 even with neg film /modern digital capable of much more is because that's what looks good. HDR technology now means you can reproduce your harsh midday sunlit scene perfectly, and it will look identically awful compared to the original.

The debate boils down to this: Does an image with a 300:1 dynamic range look good because it represents a physical sweetspot -- something to do with our perceptual system that works well at that ratio? Or is it that all we've had access to for hundreds of years are reflective images with a roughly 300:1 dynamic range, so we are accustomed to that?

I had a similar question in my mind before seeing the Sunnybrook HDR display. Now, after looking at a HDR image on a 50,000:1 HDR display, I am no longer concerned about over-brightness, 50,000:1 is still way less than the brightness of looking directly at the sun. It wasn't blinding. It isn't a question of harsh. Images simply looks better when they look more real.

In the coming decade, HDR digital imaging technology will arrive, and change how we take, manipulate, store, use and display images forever.