

Dynamic exposures

Julian Bucknall explains how computer science works to increase the dynamic range of your photographs

The sensor in a modern digital camera has less dynamic range than the human eye. That's why we're often disappointed with photographs we take: we don't see the sky as washed out, or the shadows as dark as they appear in our photos. Naturally, there are now ways to circumvent this using the power of the PC: enter the high dynamic range image processing algorithms.

When we purchase a digital camera, we're often concerned with the resolution of the sensor (the number of megapixels), whether it produces images in JPG or RAW format, and whether we can use different lenses to get images from close up or far away. We're not generally concerned with the dynamic range of the sensor in the camera – in other words, the range of light levels that the sensor can capture.

It turns out that old-style film has less dynamic range than a CCD (charge-coupled device) – the sensor that registers light information that's built into modern digital cameras. If you like, we've moved forward in terms of dynamic range and also, incidentally, in terms of noise: film is noisier at low light intensities than digital. Both types of camera perceive a narrower range of light levels than the human eye. That's why when we take a photograph of a landscape, for example, we get an image that doesn't register the cloud formations in the sky (the sky becomes washed out) and the shadows become undifferentiated black. If you're anything like me, you tend to get a little frustrated and disappointed that the camera isn't recording what you're seeing exactly, but I dare say we've all become rather used to the problem.

What is HDR?

There is, however, a way around this called HDR image processing (high dynamic range). This is a set of algorithms that process images to increase their dynamic range. With HDR, it's possible to produce an image that has a much greater range of light levels to approximate what the human eye can see, or even to make fantasy images that look nothing like real life. However, I'll also note that you will run

into some issues. For example, the monitors we use to view images also have a smaller dynamic range than the human eye.

Manual adjustments

Back in the days of film cameras, it was possible to increase the dynamic range of a photo when you printed the image after developing the film. Photographers like Ansel Adams were experts at using this kind of image manipulation – known as dodging and burning – to produce the dramatic photos we've all seen and perhaps bought as posters. Dodging decreases the exposure of the print making the area lighter in tone, whereas burning increases it making the area darker in tone.

Recall that in film photography, the film is a negative version of the photo. Dark areas on the negative will show up as lighter areas on the print paper because the light-sensitive silver salts in the paper will be less exposed, and therefore appear lighter once the print is fixed. Light areas will show up darker on the print, because more light hits the silver salts.

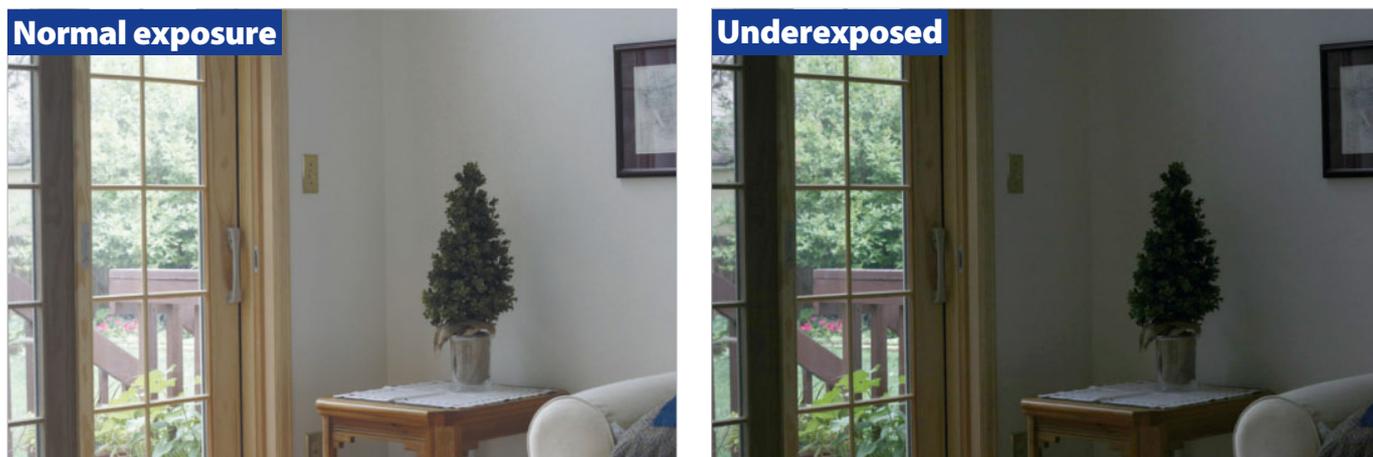
To apply dodging to the print, the photographer would cut out a shape from some opaque material like card to block off part of the scene, and then expose the print paper with that card between the projector and the paper. Because less light hit that part of the scene, it would appear lighter.

Burning was done in a similar way, but this time the photographer wanted to expose part of the scene longer than the rest. They would cut out a shape that would block off the rest of the scene, letting the part to be brightened receive more light.

There are other techniques and materials that can be used, but as you can imagine, dodging and burning this way was a labour-intensive process and was usually only done for art photos and the like. Another issue is that dodging and burning are physical manipulations that happen in real time, and it's hard to replicate the same effects across a set of prints so that the resulting images are all the same.

With digital photographs and programmatic image manipulation, it's a lot easier to create images with a higher





▲ Sample image showing normal and underexposed versions of our image. Note the variation in the outdoor colouring.

► dynamic range. The process goes like this. First you take at least three photographs of a scene. Ideally, these photos are taken with your camera on a tripod so that they all register exactly the same scene. I've tried using a hand-held camera and it just doesn't work as well during the HDR processing – there's always some obvious scene shake that can be seen. Similarly, the scene itself should be as static as possible: any moving parts (like leaves fluttering in the wind, waves crashing on the beach, or cars or people passing by, for example) won't be the same in each image, causing scene shake in the processed HDR file.

Although the photos are of the same scene and have the same focus and aperture settings, you take them using different exposure times. For best results, you should shoot one photo as normal, and the others two stops either side. Increasing exposure by a single stop doubles your exposure time, and decreasing by a stop halves it. The dynamic range in photography is the number of stops between the darkest part of an image where you can still resolve detail and the lightest part. DSLRs generally have about 11 stops of dynamic range at low ISO values, and point-and-shoot cameras a stop or so less. The external Viewsonic monitor I'm

using with my MacBook Pro has almost 10 stops, which means that the photos I take with my camera already have twice the dynamic range that my screen can show.

Some cameras, especially DSLRs, have a mode whereby you can shoot three photos as a set, the other two bracketing the first in terms of exposure. The different exposures are regulated by the camera automatically. On my Canon Rebel XTi (also known as the 400D), this is AEB mode (auto-exposure bracketing) and I can set the required +/- 2 stops there.

If your camera shoots RAW instead of JPG and your HDR image processing application supports it, it's possible to just use one photo. The results won't be nearly as good though.

Once you have your three differently exposed photos, it's time to process them. The first stage is to analyse all three photographs in order to merge them as a single HDR image.

Encoding photos

Without getting bogged down in compression details, image formats and so on, traditional images (like an individual photo from our set of three) encodes colour information for each pixel as a set of three bytes, one for red, one for green and one for blue. Each colour channel for

each pixel can therefore represent 256 different levels, and the pixel itself can be defined as a single 24-bit value.

An HDR image is different. In a simple sense, it encodes more bits per colour channel per pixel than a standard image, but that's not the entire story. To understand why, we need to understand what's meant by 'gamma correction'. Gamma defines the difference between a pixel's colour values and how bright it is actually perceived (its luminance). For a camera, if you double the amount of light on a sensor's pixel, a value twice the original is detected – the relationship between a pixel's value and its luminance is linear.

This linear relationship doesn't apply for our eyes. When we increase luminance at low light levels, we perceive a larger increase in light. At higher light levels, we don't perceive increases in luminance as well. Our eyes are more sensitive to changes in dark tones than changes in light tones. Accounting for this difference is known as gamma correction.

A camera will apply gamma correction to an image before it saves it as a JPG file. It's the same with image processing applications that work with RAW image files, when you save an image as a JPG. In other words, instead of encoding the values of the pixels reported by the sensor, it will apply a gamma correction (the industry standard value is 1/2.2) first. Instead of recording the real values of the pixel colours, it will encode them as colour tones that we perceive as varying uniformly. In effect, the code that saves an image as a JPG file uses more bits to encode darker colours than lighter ones. When you view such a gamma-encoded JPG on your screen, the screen software will apply the reverse gamma correction (a gamma of 2.2 usually) so that the image you see is roughly the same as the original scene. Lighter shades will have less variation than darker shades, though.

Notice also that merging three JPGs to produce an HDR image would necessitate that the gamma correction be reversed before the merge. JPG is already a lossy compression algorithm (we're losing image information when the camera creates a JPG, smearing out

Spotlight on... The Lytro camera

Although it's not yet been released, the Lytro light field camera promises a new way to take photos without worrying about focus.

The photographs this camera takes aren't ordinary JPGs like we get from the cameras we use today. The camera records a lot more information about the light falling on the sensor in such a way that post-processing on a PC using ray-tracing algorithms will let the user to produce a photograph at any focal length. In other words, you don't have to focus on the subject with the camera. You can 'focus' on any part of the scene later on in the supplied image processing application on your PC. Therefore, one shot, post-processed, could produce images where the focus is on

a subject in the foreground or where the focus is on a subject in the background.

Let's imagine a similar system where the camera records many exposures at one time. Later on, during the post-processing phase, we can select exposures and merge them to produce HDR images automatically. Or maybe it's done for us automatically by the software. Either way we don't have to take three photos of a scene in quick succession; the camera takes one shot and the image processing application takes care of the rest. This imagined camera won't suffer from the issues of creating a high-dynamic-range sensor – at least one of the photos will be a normal image taken at the correct exposure. ■

Over exposed**Tone mapped HDR**

▲ An overexposed version of our shot and a tone mapped version of the HDR image. HDR looks a little unrealistic, but preserves the important data.

Detail in shadows

Modern camera sensors generally have a good response to light levels across the spectrum from light to dark. Unfortunately, the darker a region of a scene, the less light reaches the sensor, and the more noise is generated (or perceived) as a consequence. You can see this if you zoom into an otherwise acceptable photograph: noise is seen as random pixels of different colours, and in a dark area those pixels become more noticeable, just because they can be brighter.

As it happens, the extra information that's produced by recording a scene at three levels of exposure – especially that at +2 stops of extra exposure – will fill in the details of the dark area, and the merging of the three exposures into a single HDR image will remove the noise. ■

high-variation regions of the image), which means that our three differently-exposed photos should all be in RAW format.

Now back to the HDR image. It converts colours as tones using the gamma correction just discussed (so darker tones have more shades) and encodes them as a floating point number, using either single precision (32-bit) or double precision (64-bit). All three standard images are used in this conversion/encoding process, since each will have a different complementary set of tone information. There's a lot of redundant information, so the most popular HDR format compresses the RGB tone data as three fixed point values, with an extra byte holding a common exponent.

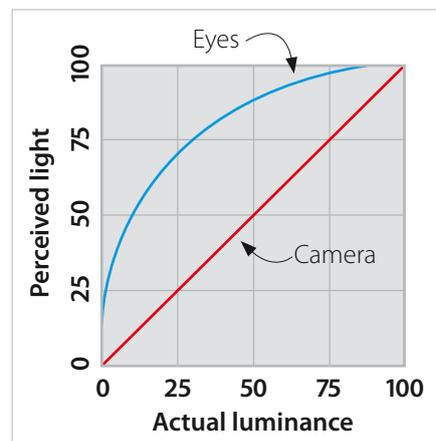
Viewing the image

So now we have an HDR image whose tones are encoded in high fidelity using a special encoding scheme. To view it, however, we have to export it as a JPG again. Because of the +/- 2 stops bracketing, the HDR image has a dynamic range of roughly 14 stops, which happens to be about as much as the eye can perceive, but we have to convert that wide dynamic range into the 10 stops of the screens we use. We have to map the HDR tones into the range that can be displayed on a monitor. Since we're going to be removing

information (or reducing the amount of data) in order to do this, tone mapping is a lossy one-way process. It also tends to be an interpretive process, much as dodging and burning was to the film photographers, requiring a good eye and attention to detail and lots of time tweaking and experimenting.

In any HDR image-processing application, there will be several knobs or sliders you can adjust to change the tone mapping, and thereby the look and feel of the finished low dynamic range image (sometimes known as LDR). Some examples include exposure or brightness, gamma or contrast – sharpening tools that alter the local appearance of parts of the image. Some may come with presets that generate an LDR image according to the image style you want to show (Photomatix, the application I use, has a preset called Grunge and yes, the results are pretty grungy).

There are several tone-mapping algorithms that can process an HDR image to produce an LDR version (a hands-off algorithm), but it's a field that is being actively researched and is changing rapidly. The biggest issue that automated algorithms have to solve is that tone mapping, by its very nature, reduces contrast (contrast is a difference in luminance: the higher the contrast, the greater the difference in luminance or tone).



▲ How the eye perceives light intensity (blue line) compared to the way the camera does (red line).

Improving the sensors

The techniques described in the main article stem from a 1995 paper by Mann and Picard, describing the method of combining three or more images of the same scene taken with different exposures.

Another way to produce such an image is to improve the dynamic range of the sensor, from the 11 stops of a high-end DSLR to the 14 or so seen by the human eye. This camera would take photos directly in HDR format; there would be no need to take three photos and then merge them.

Image file sizes would be larger, and a standard tone mapping would be needed to convert an HDR image to an LDR one the user can see on the back of the camera, and tone mappings are processor intensive. An HDR image processing PC application would have to be provided too. ■

The earliest hands-off algorithms were global algorithms that acted on the image as a whole. An example of global tone mapping is to take the tone (or luminance) information and scale it linearly to the required range. Luminance is a value on an exponential scale (each increase of one stop doubles the brightness), so the linear scaling is usually done on the logarithm of the tone values. The problem with this was that it paid no attention to areas within the image with local high contrast.

Other algorithms attempt to take into consideration regions of the image that have wide variations in contrast. The algorithms use a range of methods to generate an LDR image as a multi-pass process. Some, for example, tone map most of the image one way, then map the localised high-variation regions using another method to maintain contrast.

As you can see, the problem of HDR image processing hasn't been completely solved yet. Nevertheless, it's fast becoming an accepted way to process images with high contrast, and can produce some stunning photographs. **PCP**

Julian M Bucknall has worked for companies ranging from TurboPower to Microsoft, and is now CTO for Developer Express. He explained how to spot faked photos in issue 307. feedback@pcplus.co.uk