

**NEW
EDITION**
REVISED
& UPDATED

Practical HDR

Second Edition

A complete guide to creating High Dynamic Range images with your digital SLR

David Nightingale



Practical HDR





Practical HDR

Second Edition

The complete guide to creating High Dynamic
Range images with your digital SLR

David Nightingale



AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Focal Press is an imprint of Elsevier



Focal Press is an imprint of Elsevier
225 Wyman Street, Waltham, MA 02451, USA

Copyright © 2012 The Ilex Press.
All rights reserved
This book was conceived, designed,
and produced by Ilex Press Limited
210 High Street, Lewes, BN7 2NS, UK

Publisher: Alastair Campbell
Creative Director: Peter Bridgewater
Associate Publisher: Adam Juniper
Managing Editor: Natalia Price-Cabrera and Zara Larcombe
Editor: Tara Gallagher
Associate Editor: Steve Luck
Art Director: Julie Weir and James Hollywell
Senior Designer: Emily Harbison and Kate Haynes
Designer: Richard Wolfströme and Ginny Zeal

No part of this publication may be reproduced, stored in a
retrieval system or transmitted in any form or by any means
electronic, mechanical, photocopying, recording or otherwise
without the prior written permission of the publisher

Permissions may be sought directly from Elsevier's Science
& Technology Rights Department in Oxford, UK:
Phone (+44) (0) 1865 843830; Fax (+44) (0) 1865 853333;
Email: permissions@elsevier.com.
Alternatively visit the Science and Technology Books website
at www.elsevierdirect.com/rights for further information

Notice - No responsibility is assumed by the publisher for any
injury and/or damage to persons or property as a matter of
products liability, negligence or otherwise, or from any use
or operation of any methods, products, instructions or ideas
contained in the material herein

Library of Congress Control Number: A catalog record for
this book is available from the Library of Congress

ISBN: 978-0-240-82122-1

For information on all Focal Press publications visit our
website at:
www.focalpress.com

Trademarks/Registered Trademarks
Brand names mentioned in this book are protected by
their respective trademarks and are acknowledged

Printed and bound in China

09 10 11 12 12 11 10 9 8 7 6 5 4 3 2 1



Contents

6 Introduction

1 Understanding Dynamic Range

10 Real World Dynamic Range
12 Digital Sensors and Contrast Ratios
14 Using and Understanding
the Histogram
18 HDRI and Tone Mapping Explained

2 Shooting for HDR

24 Camera Settings and Equipment
28 Metering the Dynamic Range
of a Scene
30 Shooting a Bracketed
Exposure Sequence

3 Merging Your Bracketed Sequence

38 HDR Software: The Options
40 Pre-processing Raw Files
42 Adobe Photoshop CS5
46 Photomatix Pro
50 HDR Express
54 HDR Efex Pro
58 Oloneo PhotoEngine

4 Creating Photorealistic HDR Images

66 Photorealism and
HDR Photography
68 Photoshop's HDR Options
74 Photomatix Pro's Tone Compressor
78 Photomatix Pro's Exposure Fusion
80 HDR Express
84 HDR Efex Pro
88 Oloneo PhotoEngine
94 HDR Showcase #1

5 Creating Hyper-real HDR Images

102 Photomatix Pro:
The Details Enhancer
108 The Details Enhancer:
Worked Example
112 Photoshop Merge to HDR Pro
116 HDR Express
118 HDR Efex Pro
122 Oloneo PhotoEngine
124 PhotoEngine vs. Photoshop
126 Which Program?
128 HDR Showcase #2

6 HDR Post-production

138 Removing Noise
142 Removing Halos
146 Enhancing Contrast Using Curves
150 Creating Single Image HDRs
154 Enhancing Low Contrast Scenes
156 Merging Two HDR Images
166 HDR Showcase #3

172 Glossary
174 Index
176 Acknowledgments



Introduction

When I started out in photography, one of the things I found most frustrating was that my camera seemed to have a totally different view of the world to me. Where I saw a glorious sunset, my camera saw an overexposed wash of pale colors; where I saw clouds floating over a beautiful landscape, my camera saw a flat white sky, or a dark foreground, or both; and backlit portraits would all end up as silhouettes. With practice, I learnt a variety of techniques that improved my photography: using exposure compensation, setting my camera to Manual, and using different metering patterns to assess a scene. In other words, I developed a better understanding of how to compensate for the various ways in which my camera's perception of the world was more limited than my own.

At the same time, I also learnt that there are some shots that you just can't take—at least not without using additional lighting, reflectors, neutral density graduated filters, and so on. These are the shots where the contrast is so high that some areas of the image will end up overexposed, underexposed, or both. The problem is that while our eyes can see a full range of tones in these high contrast scenes, our digital cameras cannot.

This problem isn't a new one, and the earliest photographers looked for different ways to record high contrast, or "high dynamic range" images. In the 1850s, Gustave Le Gray produced a number of dramatic seascapes constructed from two negatives—one exposed for the sea, the other for the sky. He cut both negatives along the horizon, then used the two parts to create a single photographic print. In this way he was able to capture all of the detail in the scene,

which would have been impossible with a single exposure. Motivated by the same problem, Charles Wyckoff developed a wide dynamic range film composed of three layers, each of which had a different sensitivity to light. He used this to produce photographs of nuclear explosions, which first appeared on the cover of *Life* magazine in the 1940s.

But it wasn't until much more recently that what we now commonly refer to as High Dynamic Range, or HDR, photography began to develop. In 1985, Gregory Ward created the Radiance RGBE file format for HDR images—a format that is still in use today—while in 1993, Steve Mann reported creating a tone mapped image from a sequence of exposures of normal digital images. The idea was simple—shoot a sequence of exposures that covered the full brightness range of the image, and combine them into a single, high dynamic range picture that would contain detail in everything from the brightest highlight to the deepest, darkest shadow.

This sounds straightforward, but there are numerous issues that can make it far from easy. To start with, you need to be able to meter the scene to calculate the number of exposures to make sure to capture the entire dynamic range. You also need to know how to deal with any significant movement between the frames, and you have to understand how to create and "tone map" your images to create a final picture that meets your creative expectations. As you will see in this book, once you understand these issues, HDR imaging is a powerful technique that can be used to produce photographs that are simply not possible through any other means.





Chapter 1:
Understanding Dynamic Range

Real World Dynamic Range

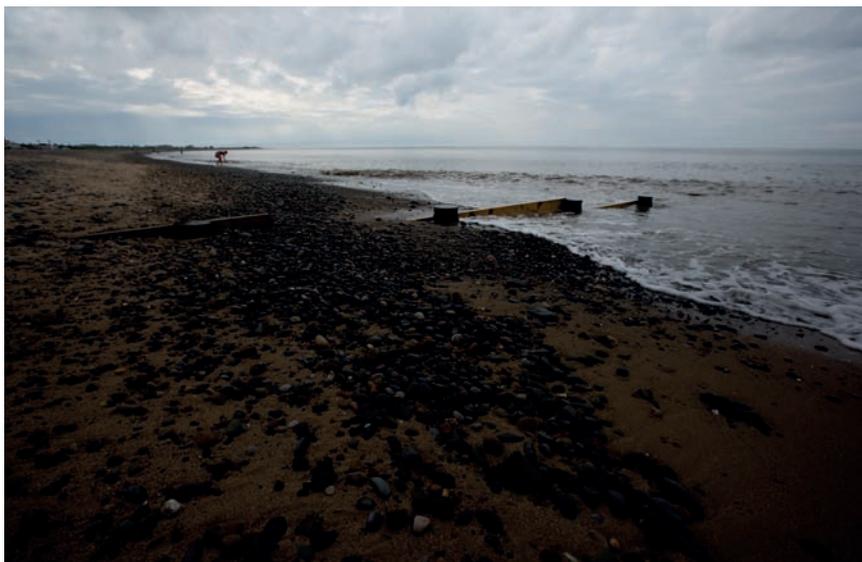
One of the first things you learn as a photographer is that the way you perceive a scene is often quite different to how your camera evaluates and processes the same data. For example, a shot you intend as a backlit portrait may well end up as a silhouette, a shot of a brightly lit scene may end up looking too dark, and so on. In other words, the image you see is sometimes not the one you manage to take.

There are two main reasons for this. The first is that your camera will attempt to set an optimum exposure for a particular shot, so it will set the aperture or shutter speed, or both, to make sure that enough light hits the sensor to produce a well-exposed image.

The problem is, your camera assumes an average level of illumination for every shot you take, so if you shoot in a dark room, the image may well end up brighter than the original scene as your camera bases its exposure on an average level of illumination. By the same token, a shot of a person, backlit by a bright sky, may produce a shot with a beautifully exposed sky, but no detail in your subject's face—again because the camera works on the assumption that all pictures are taken under “average” lighting conditions.

The second reason your images may not end up as you intend them is because of the difference between the way in which you see things and how your camera records them. While you might see a

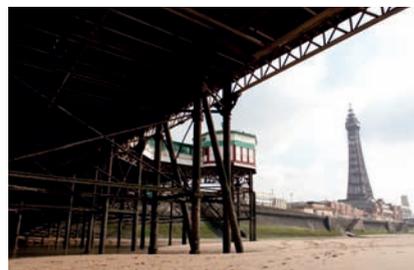
richly detailed landscape set against the backdrop of a bright, but cloudy sky, the camera will be likely to deliver a picture where the sky looks as you intended, but the foreground is too dark, or with a correctly exposed foreground and an overexposed, blank white sky. Worst of all, you might get an image with both a featureless sky and an overly dark foreground. In this instance, adjusting the exposure won't fix the problem, it will just present you with a different one. Increasing the exposure to compensate for an overly dark foreground will overexpose the sky still further, while decreasing the exposure to retain the detail in the sky will lead to an even darker foreground.



A correctly exposed sky with an overly dark foreground



A correctly exposed foreground with an overexposed sky



An image containing areas of under and overexposure



An HDR image allows you to combine the brightest and darkest parts of a scene that your camera wouldn't be able to record in a single shot.

The problem is that the “dynamic range” of the scene you are trying to photograph is larger than your camera can record. At its simplest, dynamic range is the ratio between the lightest and darkest tones in an image. This is often measured in EV (Exposure Value), which refers to the combinations of shutter speed and relative aperture that give the same exposure. Additional terms that are often used in this context are “f-stops” and “contrast ratio.”

When we view a scene, we can perceive a range of between 10–14EV (or f-stops), and if we take into account the fact that our eyes can adjust to different levels of brightness this increases to around 24EV. But digital cameras can only record an EV range of 5–9 stops, so if the scene contains an EV range of around 9 or more, there is no combination of

aperture or shutter speed that will allow you to capture the entire dynamic range of the original scene. All you can do is optimize the exposure for the shadow detail or the highlight detail, but you will inevitably lose one or the other.

From a photographic point of view, there are three solutions. First, you can simply avoid taking shots where you know that the size of the EV range will compromise the quality of the final image. In the case of a landscape shot this might mean waiting until the balance of light between the brightest and darkest areas of the scene falls within a smaller EV range.

Second, you could use a graduated neutral density (ND) filter to darken the brightest areas of the image. These can darken a portion of an image in increments from 1–3EV and are

especially useful for landscape photography as the sky can often be a lot brighter than the foreground. However, the gradation on the filter is fixed, so it is generally only suitable if there is a clear “dividing line” between the light and dark areas of the scene—useful for a landscape with a flat horizon, but less useful for images where the bright and dark areas are irregularly shaped.

Finally, you can shoot a range of different exposures that record the detail in both the deepest shadows and the brightest highlights, and then combine them into an High Dynamic Range (HDR) image, which is clearly the key topic for the remainder of this book.



JPEG files distribute the data in a slightly different way to Raw files. While there is still a bias toward the highlight detail, the number of levels allocated to each area of brightness is more evenly distributed, as illustrated in the accompanying grid.

12-bit Raw file	Levels	8-bit JPEG file	Levels
1: Highlight areas	2048	1: Highlight areas	69
2: Bright areas	1024	2: Bright areas	50
3: Midtones	512	3: Midtones	37
4: Dark areas	256	4: Dark areas	27
5: Shadow areas	128	5: Shadow areas	20

Digital Sensors and Contrast Ratios

Your eyes and your camera's sensor respond to light in a similar way, with photons striking a receptor (either your retina or the photosites on the sensor). This generates a signal, the strength of which is proportional to the amount of photons striking the receptor within a given amount of time. With the eye, exposure to light is continuous, and the effect on the retina decays over time, but with a camera the effect is additive—the shutter opens, photons are collected in the photosites for a finite period of time, and then the shutter closes. At this point, the photons are “counted” to produce a digital signal indicating the amount of light that was received during the exposure.

The reason a typical digital camera can capture an EV range of between 5 and 9 stops is because each photosite on your camera's sensor has a specific capacity, and once it reaches the maximum limit the signal it outputs is the same, irrespective of how many further photons strike it. In other words, it's like a bucket: once it's full there is no way to add any further content. This is significant, as differences in brightness can only be recorded in terms of the differences in the number of recorded photons. For example, if the capacity of each photosite was a maximum of 1024 photons, then the theoretical maximum EV range it could

record would be 10EV (210 levels). In practice, this theoretical limit can never be reached as recording detail in the very darkest areas of an image is inaccurate due to the inherent noise that sensors generate. No analog-to-digital converter is sufficiently accurate to distinguish between single photons.

Another problem is that while a 12-bit Raw file is capable of recording a theoretical maximum of 4,096 levels of brightness, the amount of data allocated to the different levels of brightness within the scene varies, with half (2,048) allocated to the brightest areas, half of the remaining 2,048 allocated to the next brightest areas, half of the remaining 1,024 to the next, and so on, as outlined in the grid at left. This means considerably more data is used to record detail within the brighter areas of an image. With a normal photograph this isn't especially significant, but with an HDR image—especially one where you really want to maximize the amount of detail in all areas of the image—the shadow detail becomes much more significant. If you haven't recorded sufficient levels of detail in the darker areas the conversion process will struggle to produce a “clean” image, resulting in noise, banding, or other digital artifacts in the darker areas of the converted image.

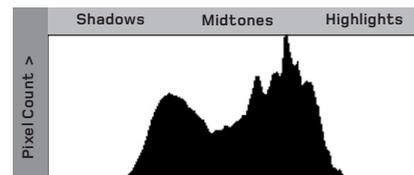
Contrast ratios of other devices

Although it would be easy to think that the solution to dealing with high dynamic range images is simply to develop cameras with sensors that can record a higher dynamic range to start with, we also need to consider how our photographs will be displayed. Our cameras are capable of recording roughly 5–9EV, which equates to a contrast ratio of 32:1–512:1, but as noted, the dynamic range of an actual scene can be quite a bit higher; the EV range of a shot containing very bright highlights and very deep shadows could be around 14EV, which is a contrast ratio of 16,384:1. Even if our camera could record this level of data, there are very few ways in which we could display the image: a photographic print typically has a contrast ratio of around 300:1, while a standard computer monitor has a contrast ratio of approximately 500:1. So even if your camera could record a much larger EV range, most standard output devices are Low Dynamic Range devices and could not display it. There are a number of High Dynamic Range monitors that have a claimed contrast ratio of around 200,000:1, but these are incredibly expensive, and unlikely to be available to the domestic market any time soon.

Using and Understanding the Histogram

Understanding the histogram is essential to every stage of producing an HDR image. You will need to use your in-camera histogram to calculate the EV range you need to shoot; to evaluate the individual shots in a bracketed exposure sequence; to change a variety of settings when you convert your HDR images; and to judge images during any subsequent post-production. In short, it's an important tool that can provide invaluable feedback as you work your way through every stage of the HDR process.

We've already examined the relationship between the EV range of a scene and the dynamic range of your camera's sensor and, as noted, the former can often be larger than the latter, resulting in areas of over- or underexposure in your images. A related term is "tonal range," which basically refers to the range and distribution of the tones between the lightest and darkest areas of an image. For example, an image with a wide tonal range will include both dark and light areas (and a range of tones in between), whereas an image with a narrow tonal range will be predominantly composed of midtones. The range and distribution of tones is graphically represented using a histogram, which is a graph with two axes. The horizontal axis shows the distribution of tones from black (at the left) to white (at the right), while the height of the graph indicates the number of pixels at any given tonal value, so the higher the peak on a histogram, the more pixels are within the image at this point.



In this example you can see that the data is clumped toward the middle of the histogram. This indicates two things. First, it shows that the tonal range is reasonably narrow—there are no blacks or whites, so no deep shadows or bright highlights. In other

words, the majority of the tones cover the midtone values. Second, it tells us that the EV range of the original scene was smaller than the dynamic range of the camera's sensor, as all the data within the original scene "fits" within the horizontal axis.



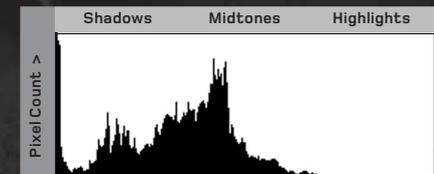
The entire EV range of the original scene has also been captured in this shot, but the shape of the histogram is very different to the first one. There is a large clump of data to the right of the histogram, and a much smaller

clump toward the left. The large peak on the right represents the wall, while the smaller clump represents the darker areas of the phone on the red wooden board.



In this example, the histogram is butted up to the right edge. This indicates that a portion of the picture is over-exposed, and if you take a look at the original image you can see why. I set the exposure to record a good range of tones

in the foreground, but because the EV range of the original scene was quite high, this caused the brightest areas of the image to be “clipped” from the histogram, meaning they ended up as pure white in the photograph.

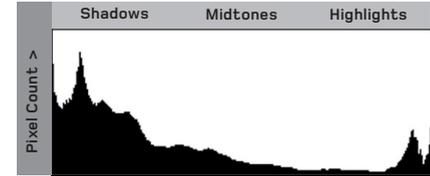


This image contains quite large areas of underexposed data, and the shadow and foreground detail have been lost. In this instance this is indicated by “clipping” at the left edge of the histogram, where

a range of areas within the original capture are black. Increasing the exposure would shift the histogram to the right, and recover some shadow detail, but perhaps not all of it.

What all these examples demonstrate is that the histogram can provide invaluable data regarding your initial exposures and, in the case of your in-camera histogram, is an instant way to accurately assess both the EV range of the scene you are photographing and the tonal range your camera has captured.

By accurately assessing your histogram you can make sure your final HDR image contains a full tonal range.



In this example, both the shadow and highlight areas have been clipped, meaning there are areas of pure black and pure white in the image. This means the EV range of the original scene is considerably larger than the

dynamic range of the camera's sensor. Adjusting the exposure to record more highlight detail would make the shadows worse, and vice versa.

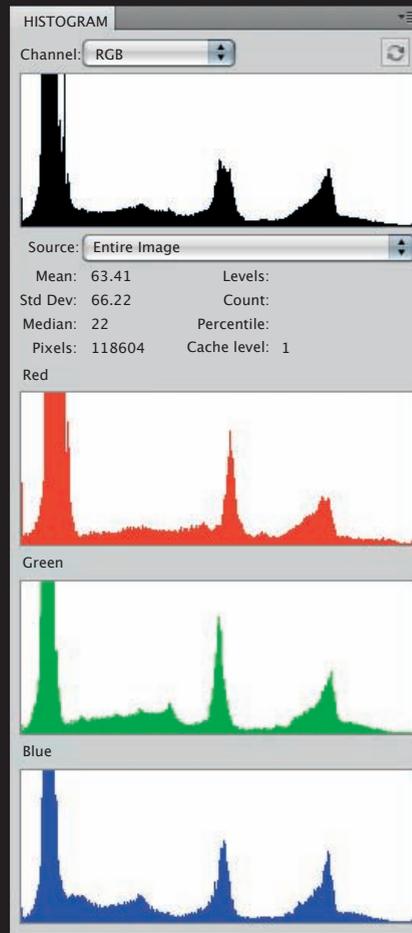


RGB and Brightness histograms

The histograms in the previous examples are all brightness histograms, which read the scene as a whole. However, some digital cameras allow you to switch to an RGB histogram that shows the three color channels (red, green, and blue) individually. An RGB histogram will often be more accurate as, while the brightness histogram might indicate that your exposure contains no shadow or highlight clipping, the RGB histogram might indicate that the highlight exposure for the green and blue channels is OK, but that the data in the red channel is overexposed, for example. Avoiding over- and underexposure is definitely one of the things you should aim for when shooting sequences of images to process into an HDR image and anything that can aid you in avoiding exposure errors is something you should embrace.



An in-camera
RGB histogram

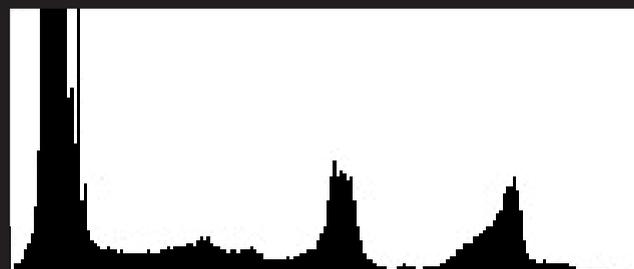


Histograms in Photoshop

When working in Photoshop, you can also assess your images using a histogram, by choosing *Window>Histogram* to open the Histogram Palette. The default histogram shows a brightness histogram, similar to the in-camera versions we just looked at, but you can also call up RGB histograms for your images.

To switch the display to show the RGB histogram, click the dropdown menu at the top-right of the histogram palette and select *All Channels View*. This will display the three channel histograms, as well as the original composite display. This can be especially useful during post-production to prevent you inadvertently clipping the detail in either the shadows or highlights of an individual channel. To make sure this doesn't happen, switch your camera to the RGB histogram as often as possible, and make sure that you get into the habit of routinely previewing your exposures as you shoot them. If you notice any under- or overexposure in one or more of the three channels, adjust your exposure and take the shot again.

Photoshop's
"All Channels View"
histogram



Photoshop's default
brightness histogram

HDRI and Tone Mapping Explained

When the EV range of a scene is larger than the dynamic range of the camera you have no option other than to accept that some of the detail cannot be recorded in a single image, and you will have to sacrifice shadow detail, highlight detail, or both. In many cases, this isn't a problem—simply set your exposure to make sure that the areas of the scene you are most interested in are exposed adequately. In this example, I set my exposure to capture the sunset and the bright reflections in the water, allowing the young woman in the foreground to become a silhouette. This works well, and the loss of detail creates a more subtle, enigmatic image.



Sunset in Dubai

Let's look at another example, this time a view underneath a pier. Here, the EV range of the original scene is about 5EV (or stops) larger than my camera's dynamic range, so a sequence of images from -2EV to +3EV is required to successfully capture the entire dynamic range. The histograms of the lightest and darkest images show the sequence as a whole covers the entire EV range of the scene, with no clipped highlights in the darkest exposure, and no clipped shadows in the lightest.

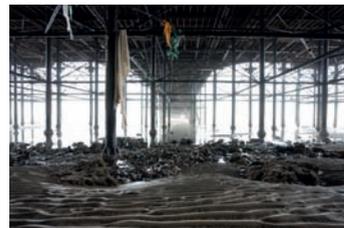
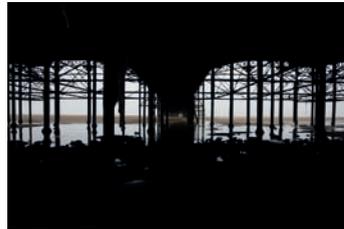
Having done this, the next step is to combine, or convert, this sequence into a single High Dynamic Range (HDR) image. The term "High Dynamic Range"—or "High Dynamic Range Image" (HDRI)—is often used to label a particular style of picture, but, technically, this name is incorrect. As an HDR image is a 32-bit image it cannot be displayed or printed on conventional low dynamic range (LDR) devices such as computer monitors or printers. In other words, most—if not all—of the HDR images you have seen (including those in this book) are actually LDR images, albeit created using a variety of HDR techniques.

Indeed, HDR images themselves are visually quite dull, at least when displayed on a standard monitor or the printed page. This is because we can only display a portion of the full dynamic range, so a "pure" 32-bit HDR file of the pier doesn't really look any better than any of the original images in the exposure sequence, and still displays a substantial amount of clipping in both the shadow and highlight areas.

The histogram for the darkest exposure (-2EV)



-2EV



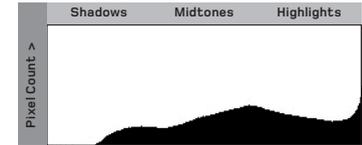
+1EV

-1EV



+2EV

The histogram for the lightest exposure (+3EV)



Metered exposure



+3EV

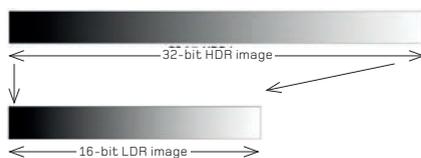
So, while an HDR image may contain all the data from the original scene—in both the shadow and highlight areas—we can't see all of the detail if we are viewing it on a conventional LDR device. This means creating an HDR file is only a part of the story, and this story isn't complete until we convert it back into an LDR image using a process called "tone mapping."



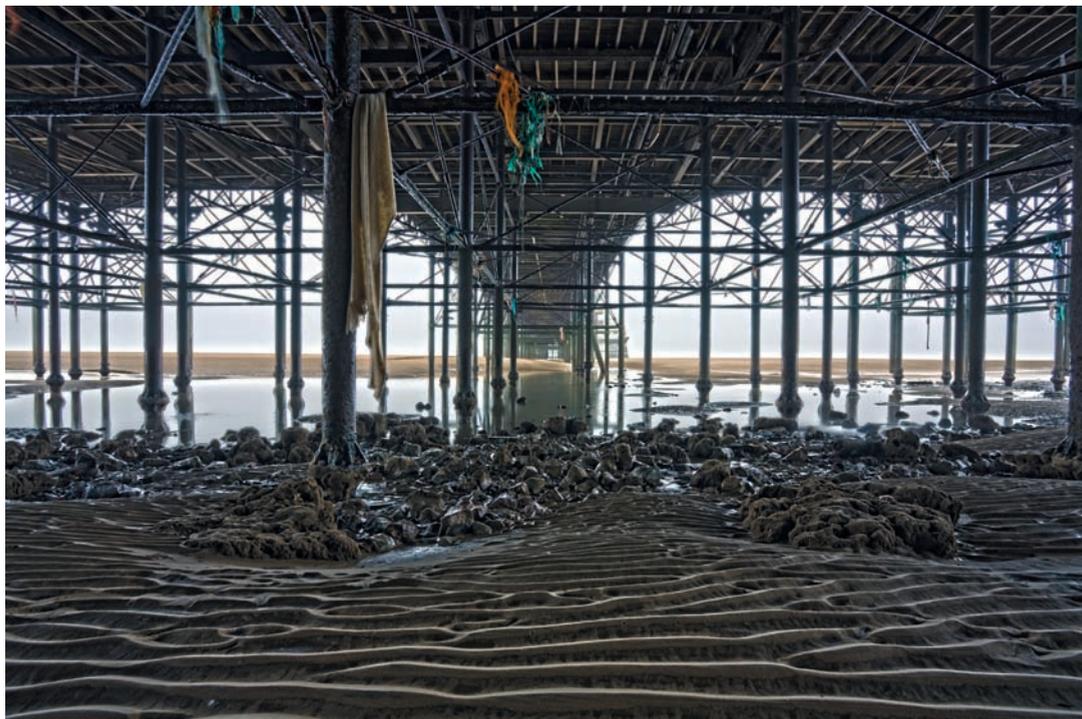
32-bit HDR file



Converted using Photoshop Highlight Compression



Tone mapping
a 32-bit HDR file



Converted using Photoshop Local Adaptation

Tone mapping

At its most basic, tonemapping a 32-bit HDR image simply involves compressing the image data into a range that can be printed or displayed. If we use one of Photoshop's basic tone mapping methods, Highlight Compression, (as discussed in chapter four) we can compress this HDR picture of the pier into a 16-bit (or 8-bit) file that contains a full range of tones in both the shadow and highlight areas (above right).

As you can see though, converting the image in this way still doesn't produce an especially good result. While we do now have a full range of tones in both

the shadow and highlight areas of the image, the underside of the pier and the foreground are both still dark, and the sky and brighter areas of the image lack any detail.

This is because this particular form of tone mapping applies a global algorithm to the 32-bit HDR image, so every pixel is converted—or “mapped”—in the same way. Essentially, the luminance of each pixel is altered in a uniform way to make sure that all the data fits into the smaller luminance range of the LDR image.

While these global forms of tone mapping can produce good results when the variance in luminance is

evenly distributed, a form of “local” tone mapping will often produce a better result. Local tone mapping adjusts each pixel in relation to the value of the surrounding pixels in the image, rather than treating them all the same. The math behind this process is extremely complex, but the result is an increase to the local contrast in different areas of an image by unequal amounts. As the final version of the pier image reveals (above left), using the localized tone mapping algorithms in Photoshop's Local Adaptation mode has produced a considerably more striking picture (above).



A tone mapped HDR image



The original image



Increased local contrast

In this Local Adaptation example, there is now a high degree of contrast in the foreground of the image and the underside of the pier—areas that were quite dark when converted using Photoshop's simpler, global Highlight Compression method. There is also much greater detail and contrast in the brighter areas of the image. Put simply, using a local tone mapping algorithm has produced a much better image.

Throughout this book, the majority of the techniques will concentrate on local tone mapping algorithms to produce an LDR image, though the extent to which this effect is applied will vary, depending on the result we are looking for.

The image of the shopping cart and trash can (above), looks more like an illustration than a photograph once the HDR file has been converted into an LDR image. The final image was produced from a single Raw file (a technique covered in chapter six) and was processed to maximize the local contrast. In the original shot, the trash can is very dark, but in the tone mapped version it shares the same tonal range as other sections of the image. In other words, the naturally occurring global contrast in the original—where some areas of the image are clearly darker than others—has been all but removed in the final, tone mapped image.

By comparison, the tone mapped image of a building and cranes—which required seven bracketed exposures to capture the entire dynamic range—looks far more natural. Although the dynamic range of the original scene was very high, less local contrast was introduced at the tonemapping stage, allowing the darker areas to remain quite dark and the lighter areas to retain a higher amount of brightness. As a result, the final image appears more photorealistic than the shopping cart, although both images would still be labeled “HDR.”

