Introduction

01

There are many applications that involve digital images. They are created with modern digital cameras and scanners, rendered with advanced computer graphics techniques, or produced with drawing programs. These days, most applications rely on graphical representations of some type.

During their lifetime, digital images undergo a number of transformations. First, they are created using one of the previously cited techniques. Then they are stored on a digital medium, possibly edited via an image-processing technique, and ultimately displayed on a computer monitor or printed as hardcopy.

Currently, there is a trend toward producing and using higher-resolution images. For example, at the time of writing there exist consumer-level digital cameras that routinely boast 5- to 6-megapixel sensors, with 8- to 11-megapixel sensors available. Digital scanning backs routinely offer resolutions that are substantially higher. There is no reason to believe that the drive for higher-resolution images will abate anytime soon. For illustrative purposes, the effect of various image resolutions on the visual quality of an image is shown in Figure 1.1.

Although the trend toward higher-resolution images is apparent, we are at the dawn of a major shift in thinking about digital images, which pertains to the range of values each pixel may represent. Currently, the vast majority of color images is represented with a byte per pixel for each of the red, green, and blue channels. With three bytes per pixel, more than 1.6 million different colors can be assigned to each pixel. This is known in many software packages as "millions of colors."

This may seem to be an impressively large number at first, but it should be noted that there are still only 256 values for each of the red, green, and blue components of each pixel. Having just 256 values per color channel is inadequate for

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 2

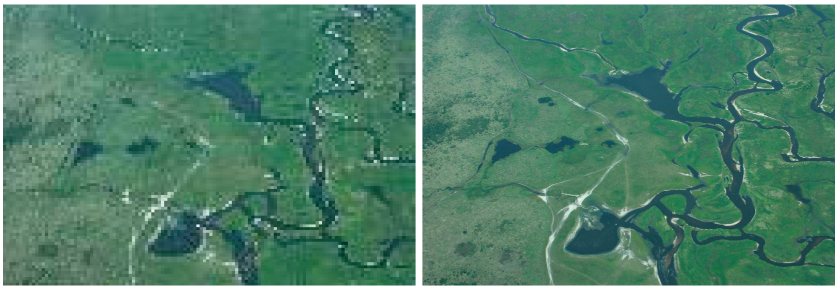

FIGURE 1.1 *Increasing the number of pixels in an image reduces aliasing artifacts. The image on the left has a resolution of 128 by 96 pixels, whereas the image on the right has a resolution of 1,024 by 700 pixels.*

representing many scenes. An example is shown in Figure 1.2, which includes an automatically exposed 8-bit image on the left. Although the subject matter may be unusual, the general configuration of an indoor scene with a window is quite common. This leads to both bright and dark areas in the same scene. As a result, in Figure 1.2 the lake shown in the background is overexposed.

The same figure shows on the right an example that was created, stored, and prepared for printing with techniques discussed in this book. In other words, it is a high-dynamic-range (HDR) image before the final display step was applied. Here, the exposure of both the indoor and outdoor areas has improved. Although this

FIGURE 1.2 *Optimally exposed conventional images (left) versus images created with techniques described in this book (right).*

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 4

FIGURE 1.3 *A conventional image is shown on the left, and an HDR version is shown on the right. The right-hand image was prepared for display with techniques discussed in Section 7.2.7.*

image shows more detail in both the dark and bright areas, this is despite the fact that this image is shown on paper so that the range of values seen is not higher than in a conventional image. Thus, even in the absence of a display device capable of displaying them, there are advantages to using HDR images. The difference between the two images in Figure 1.2 would be significantly greater if the two were displayed on one of the display devices discussed in Chapter 5.

A second example is shown in Figure 1.3. The image on the left is a conventional image shot under fairly dark lighting conditions, with only natural daylight being available. The same scene was photographed in HDR, and then prepared for display

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 5

FIGURE 1.4 *The image on the left is represented with a bit depth of 4 bits. The image on the right is represented with 8 bits per color channel.*

with techniques discussed in this book. The result is significantly more flattering, while at the same time more details are visible.

The range of values afforded by a conventional image is about two orders of magnitude, stored as a byte for each of the red, green, and blue channels per pixel. It is not possible to directly print images with a much higher dynamic range. Thus, to simulate the effect of reducing an HDR image to within a displayable range, we reduce a conventional photograph in dynamic range to well below two orders of magnitude. As an example, Figure 1.4 shows a low-dynamic-range (LDR) image (8 bits per color channel per pixel), and the same image reduced to only 4 bits per

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 6

Condition	Illumination (in cd/m^2)
Starlight	10^{-3}
Moonlight	10^{-1}
Indoor lighting	10^{2}
Sunlight	10^{5}
Max. intensity of common CRT monitors	10^{2}

TABLE 1.1 *Ambient luminance levels for some common lighting environments (from Wandell's book* Foundations of Vision [135]).

color channel per pixel. Thus, fewer bits means a lower visual quality. Although for some scenes 8 bits per color channel is enough, there are countless situations in which 8 bits is not enough.

One of the reasons for this is that the real world produces a much greater range than the two orders of magnitude common in current digital imaging. For instance, the sun at noon may be 100 million times brighter than starlight [34,120]. Typical ambient luminance levels for commonly encountered scenes are outlined in Table 1.1.[1]

The human visual system is capable of adapting to lighting conditions that vary by nearly 10 orders of magnitude [34]. Within a scene, the human visual system functions over a range of about five orders of magnitude simultaneously.

This is in stark contrast to typical CRT (cathode-ray tube) displays, which are capable of reproducing about two orders of magnitude of intensity variation. Their limitation lies in the fact that phosphors cannot be excited beyond a given limit. For this reason, 8-bit digital-to-analog (D/A) converters are traditionally sufficient

1 Luminance, defined in the following chapter, is a measure of how bright a scene appears.

for generating analog display signals. Higher bit depths are usually not employed, because the display would not be able to reproduce such images at levels that are practical for human viewing.[2]

A similar story holds for typical modern liquid crystal displays (LCD). Their operating range is limited by the strength of the backlight. Although LCD displays tend to be somewhat brighter than CRT displays, their brightness is not orders of magnitude greater.

In that current display devices are not capable of reproducing a range of luminances anywhere near the capability of the human visual system, images are typically encoded with a byte per color channel per pixel. This encoding normally happens when the image is captured. This situation is less than optimal because much of the information available in a scene is irretrievably lost at capture time.

A preferable approach is to capture the scene with a range of intensities and level of quantization representative of the scene, rather than matched to any display device. Alternatively, images should at a minimum contain a range of values matched to the limits of human vision. All relevant information may then be retained until the moment the image needs to be displayed on a display device that cannot reproduce this range of intensities. This includes current CRT, LCD, and plasma devices, as well as all printed media.

Images that store a depiction of the scene in a range of intensities commensurate with the scene are what we call HDR, or "radiance maps." On the other hand, we call images suitable for display with current display technology LDR.

This book is specifically about HDR images. These images are not inherently different from LDR images, but there are many implications regarding the creation, storage, use, and display of such images. There are also many opportunities for creative use of HDR images that would otherwise be beyond our reach.

Just as there are clear advantages to using high-image resolutions, there are major advantages in employing HDR data. HDR images and video are matched to the scenes they depict, rather than the display devices they are meant to be displayed on. As a result, the fidelity of HDR images is much higher than with conventional imagery. This benefits most image processing that may be applied during the lifetime of an image.

2 It would be possible to reproduce a much larger set of values on CRT displays at levels too low for humans to perceive.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 8

FIGURE 1.5 *Color manipulation achieved on an HDR capture (left) produced the image on the right. The left-hand HDR image was captured under normal daylight (overcast sky). The right-hand image shows a color transformation achieved with the algorithm detailed in Section 7.2.7.*

As an example, correcting the white balance of an LDR image may be difficult due to the presence of overexposed pixels, a problem that exists to a lesser extent with properly captured HDR images. This important issue, which involves an adjustment of the relative contribution of the red, green, and blue components, is discussed in Section 2.6. HDR imaging also allows creative color manipulation and better captures highly saturated colors, as shown in Figure 1.5. It is also less important to carefully light the scene with light coming from behind the photographer, as demonstrated in Figure 1.6. Other image postprocessing tasks that become easier with the use of HDR data include color, contrast, and brightness adjustments.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 9

FIGURE 1.6 *Photographing an object against a bright light source such as the sky is easier with HDR imaging. The left-hand image shows a conventional photograph, whereas the right-hand image was created using HDR techniques.*

Such tasks may scale pixel values nonlinearly such that parts of the range of values require a higher precision than can be accommodated by traditional 8-bit pixel encodings. An HDR image representation would reduce precision errors to below humanly detectable levels.

In addition, if light in a scene can be accurately represented with an HDR image, such images may be effectively used in rendering applications. In particular, HDR images may be used as complex light sources that light conventionally modeled 3D geometry. The lighting effects thus obtained would be extremely difficult to model in any other way. This application is discussed in detail in Chapter 9.

Further, there is a trend toward better display devices. The first prototypes of HDR display devices have been around for at least two years at the time of writing [114, 115] (an example is shown in Figure 1.7). Their availability will create a much larger market for HDR imaging in general. In that LDR images will look no better on HDR display devices than they do on conventional display devices, there will be an increasing demand for technology that can capture, store, and manipulate HDR data directly.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 10

FIGURE 1.7 *Sunnybrook Technologies prototype HDR display device.*

It is entirely possible to prepare an HDR image for display on an LDR display device, as shown in Figure 1.2, but it is not possible to reconstruct a high-fidelity HDR image from quantized LDR data. It is therefore only common sense to create and store imagery in an HDR format, even if HDR display devices are ultimately not used to display it. Such considerations (should) play an important role, for instance, in the design of digital heritage and cultural archival systems.

When properly displayed on an HDR display device, HDR images and video simply look gorgeous! The difference between HDR display and conventional imaging is easily as big a step forward as the transition from black-and-white to color tele-

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 11

vision. For this reason alone, HDR imaging will become the norm rather than the exception, and it was certainly one of the reasons for writing this book.

However, the technology required to create, store, manipulate, and display HDR images is only just emerging. There is already a substantial body of research available on HDR imaging, which we collect and catalog in this book. The following major areas are addressed in this book.

Light and color: HDR imaging borrows ideas from several fields that study light and color. The following chapter reviews several concepts from radiometry, photometry, and color appearance and forms the background for the remainder of the book.

HDR image capture: HDR images may be created in two fundamentally different ways. The first method employs rendering algorithms and other computer graphics techniques. Chapter 9 outlines an application in which HDR imagery is used in a rendering context.

The second method employs conventional (LDR) photo cameras to capture HDR data. This may be achieved by photographing a static scene multiple times, varying the exposure time for each frame. This leads to a sequence of images that may be combined into a single HDR image. An example is shown in Figure 1.8, and this technique is explained in detail in Chapter 4.

This approach generally requires the subject matter to remain still between shots, and toward this end the camera should be placed on a tripod. This limits, however, the range of photographs that may be taken. Fortunately, several techniques exist that align images, remove ghosts, and reduce the effect of lens flare, thus expanding the range of HDR photographs that may be created. These techniques are discussed in Chapter 4.

In addition, photo, film, and video cameras will in due course become available that will be capable of directly capturing HDR data. As an example, the FilmStream Viper is a digital camera that captures HDR data directly. Although an impressive system, its main drawback is that it produces raw image data at such a phenomenal rate that hard drive storage tends to fill up rather quickly. This is perhaps less of a problem in the studio, where bulky storage facilities may be available, but the use of such a camera on location is restricted by the limited capacity of portable hard drives. It directly highlights the need for effi-

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 12

FIGURE 1.8 *Multiple exposures (shown on the right) may be combined into one HDR image (left).*

cient file formats for storing HDR video. Storage issues are discussed further in Chapter 3.

HDR security cameras, such as the SMaL camera, are also now available (Figure 1.9). The main argument for using HDR capturing techniques for security applications is that typical locations are entrances to buildings. Conventional video cameras are typically not capable of faithfully capturing the interior of a building at the same time the exterior is monitored through the window. An HDR camera would be able to simultaneously record indoor and outdoor activities.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 13

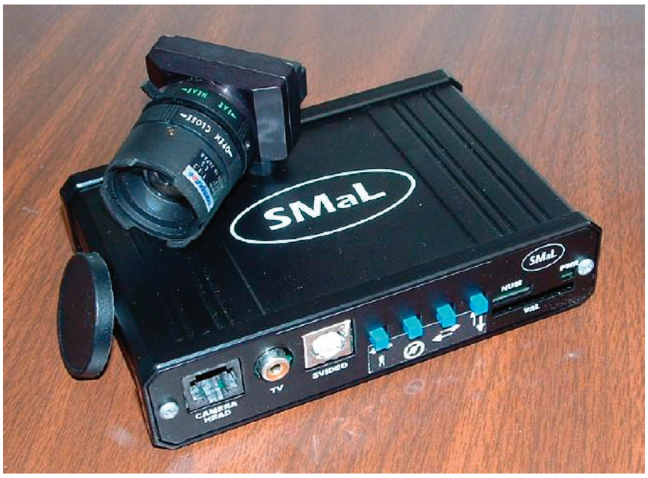

FIGURE 1.9 *SMaL prototype security camera. Image used by permission from Cypress Semicondutor.*

Many consumer-level photo cameras are equipped with 10- or 12-bit A/D converters and make this extra resolution available through proprietary RAW[3] formats (see Chapter 3). However, 10 to 12 bits of linear data affords about the same precision as an 8-bit gamma-compressed format, and may therefore still be considered LDR.

HDR image representation: Once HDR data is acquired, it needs to be stored in some fashion. There are currently a few different HDR file formats emerging. The design considerations for HDR file formats include the size of the resulting files, the total range that may be represented (i.e., the ratio between the largest repre-

3 RAW image formats are manufacturer's and often model-specific file formats containing improcessed sensor output.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 14

sentable number and the smallest), and the smallest step size between successive values. These trade-offs are discussed in Chapter 3, which also introduces standards for HDR image storage.

HDR display devices: Just as display devices have driven the use of 8-bit image processing for the last 20 years, the advent of HDR display devices will impact the general acceptance of HDR imaging technology.

Current proposals for HDR display devices typically employ a form of back-projection, either in a system that resembles a slide viewer or in technology that replaces the single backlight of an LCD display with a low-resolution but HDR projective system. The latter technology thus provides HDR display capabilities by means of a projector or LED array that lights the LCD display from behind with a spatially varying light pattern [114,115]. The display augments this projected image with a high-resolution but LDR LCD. These emerging display technologies are presented in Chapter 5.

Image-based lighting: In Chapter 9 we explore in detail one particular application of HDR imaging; namely, image-based lighting. Computer graphics is generally concerned with the creation of images by means of simulating how light bounces through a scene [42,116,144]. In many cases, geometric primitives such as points, triangles, polygons, and splines are used to model a scene. These are then annotated with material specifications, which describe how light interacts with these surfaces. In addition, light sources need to be specified to determine how the scene is lit. All of this information is then fed to a rendering algorithm that simulates light and produces an image. Well-known examples are represented in films such as the *Shrek* and *Toy Story* series.

A recent development in rendering realistic scenes takes images as primitives. Traditionally, images are used as textures to describe how a surface varies over space. As surface reflectance ranges between 1 and 99% of all incoming light, the ability of a diffuse surface to reflect light is inherently LDR. It is therefore perfectly acceptable to use LDR images to describe things such as wood grain on tables or the pattern of reflectance of a gravel path. On the other hand, surfaces that reflect light specularly may cause highlights that have nearly the same luminance as the light sources they reflect. In such cases, materials need to be represented with a much higher precision.

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 15

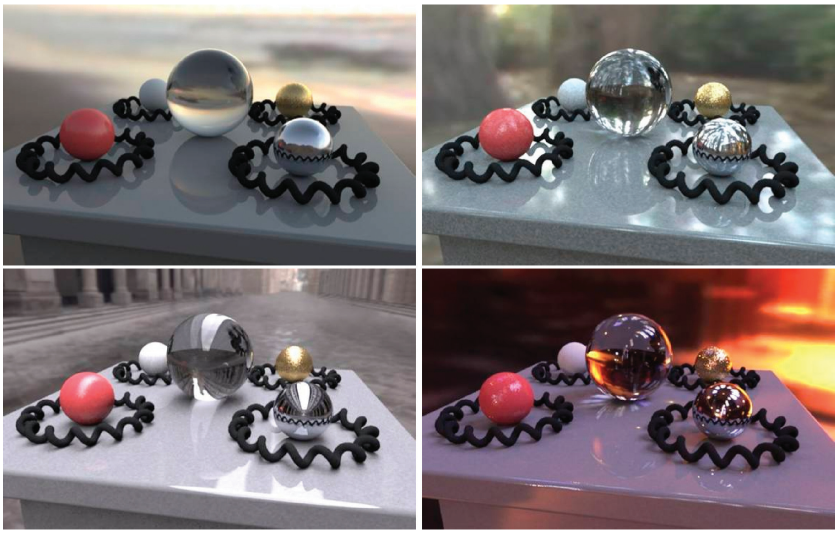

FIGURE 1.10 *HDR images may be used to light an artificial scene.*

In addition, images may be used as complex sources of light within otherwise conventional rendering algorithms [17], as shown in Figure 1.10. Here, we cannot get away with using LDR data because the range of light emitted by various parts of a scene is much greater than the two orders of magnitude available with conventional imaging. If we were to light an artificial scene with a representation of a real scene, we would have to resort to capturing this real scene in HDR. This example of HDR image usage is described in detail in Chapter 9.

Dynamic range reduction: Although HDR display technology will become generally available in the near future, it will take time before most users have made the transition. At the same time, printed media will never become HDR because this would entail the invention of light-emitting paper. As a result, there will always be a need to prepare HDR imagery for display on LDR devices.

It is generally recognized that linear scaling followed by quantization to 8 bits per channel per pixel will produce a displayable image that looks nothing like the original scene. It is therefore important to somehow preserve key qualities of HDR images when preparing them for display. The process of reducing the range of values in an HDR image such that the result becomes displayable in some meaningful way is called dynamic range reduction. Specific algorithms that achieve dynamic range reduction are referred to as tone-mapping, or tone-

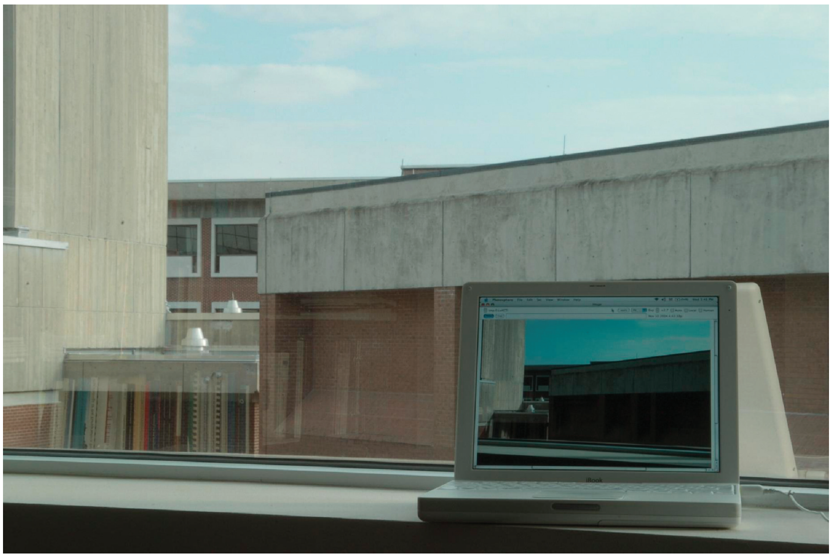

FIGURE 1.11 *The monitor displays a tone-mapped HDR image depicting the background. Although the monitor is significantly less bright than the scene itself, a good tone reproduction operator would cause the scene and the displayed image to appear the same to a human observer.*

reinhard v.2005/03/22 Prn:13/06/2005; 15:20 F:reinhard01.tex; VTEX/JOL p. 17

reproduction, operators. The display of a tone-mapped image should perceptually match the depicted scene (Figure 1.11).

In that dynamic range reduction requires preservation of certain scene characteristics, it is important to study how humans perceive scenes and images. Many tone-reproduction algorithms rely wholly or in part on some insights of human vision, not least of which is the fact that the human visual system solves a similar dynamic range reduction problem in a seemingly effortless manner. We survey current knowledge of the human visual system as it applies to HDR imaging, and in particular to dynamic range reduction, in Chapter 6.

Tone reproduction: Although there are many algorithms capable of mapping HDR images to an LDR display device, there are only a handful of fundamentally different classes of algorithms. Chapters 7 and 8 present an overview of all currently known algorithms, classify them into one of four classes, and discuss their advantages and disadvantages. Many sequences of images that show how parameter settings affect image appearance for each operator are included in these chapters.

Although the concept of HDR imaging is straightforward (i.e., representing scenes with values commensurate with real-world light levels), the implications to all aspects of imaging are profound. In this book, opportunities and challenges with respect to HDR image acquisition, storage, processing, and display are cataloged in the hope that this contributes to the general acceptance of this exciting emerging technology.