Chapter 1

Introduction

1.1 Applications Requiring a Higher Dynamic Range

In industrial processes, illumination is generally a well-controlled element of a system, so that the luminance of a scene matches well with the operation range of a camera. Industrial system developers and integrators have access to a variety of light source types, diffusers, filters, reflectors, and mirrors to design the required light pattern and uniformity. The light intensity can also be easily controlled, and light flashes or pulses can be synchronized with the camera. As the speed of the image sensors and camera increases with each new generation, it is now convenient to take multiple pictures of the same scene with light from different wavelengths or with different directions or intensities. There are, however, some applications where luminance of a scene is not completely (or not at all) under control of the user. These applications are numerous, but most of the time they do not require high dynamic range (HDR) solutions. Some applications, however, are impossible without an HDR imaging system.

HDR applications can be organized into several categories, depending on the following criteria:

- Is there relative movement between the scene and the camera?
- Is the application targeting display or machine vision?
- How much is the required dynamic range?
- Is it a real-time application or not?
- Is dynamic range extension required on the dark side, the bright side, or both?

Welding is a typical case where the main object of the scene is extremely bright (see Fig. 1.1). Exposing the sensor to this bright spot makes the image extremely dark, giving no details except, maybe, within the light arc itself. Exposing the sensor to the largest area of the scene can provide usable information, but the light arc will be totally saturated with no information. In this scene, the area of interest is the welding point, which is located very close to the brightest area, but the point itself is darker. This application is one of

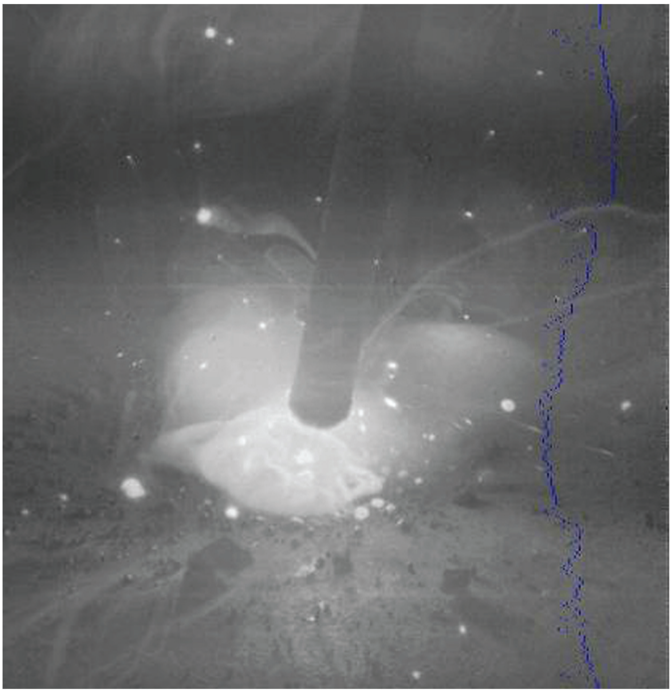

Figure 1.1 140-dB image of a welding application using a logarithmic sensor and off-chip fixed-pattern noise correction. The blue dots are an overlay of histogram data into the image (courtesy of Dierickx [83]).

the most difficult situations. The testing of explosives and electrical fuses results in similar scenes and imaging difficulties. The word "information" is used here because it indicates what is of importance in an image. We do not care much about how nice the picture looks but about how much information can be collected, processed, or analyzed. Sometimes we can accept some saturation in a picture if the saturated areas do not contain much useful information but the saturation improves the image quality around and on some objects of interest. As an example, consider an automotive camera pointed ahead of the vehicle; it is interested in the road signs, road marking, and other vehicles or pedestrians, but not the clouds in the sky.

A similar situation is encountered in automotive on-board applications such as park assist, pedestrian detection, lane departure warning, or traffic sign recognition. In these circumstances, there is no control of the general lighting that can come from various sources such as a low sun at sunset, headlights of oncoming vehicles, or sun reflection on a damp or wet road (see Fig. 1.2). These situations can become more dangerous when the road signs to be imaged are very reflective and illuminated directly by the lights of one's own vehicle, as is often the case in many countries, for example, in France.

Figure 1.2 Example of an HDR automotive front-vision application. The image was acquired in 2004 with an early-model panoramic VGA global shutter automotive image sensor.

Intelligent airbag systems are now measuring the size of the person sitting in the passenger seat to differentiate a child from an adult or a large bag, so that proper airbag deployment is achieved. Such systems make use of synchronized pulsed infrared light. This small amount of infrared light is the signal to be detected; large amounts of sunlight have greater intensity by

Figure 1.3 Image of a lightbulb using a logarithmic sensor with off-chip fixed-pattern noise correction (courtesy of Dierickx [83]).

orders of magnitude. Separating the useful part of the light from the sunlight and other light sources requires HDR techniques.

In road traffic monitoring, the situation is similar to on-board automotive applications, as scene illumination is not under control, and the oncoming vehicles can be oriented with their lights facing the camera. This situation is usually solved in traffic monitoring by placing the camera higher, so that the lights do not shine directly into it. However, when working with tunnel entrances and exits, bright light from the sun outside the tunnel contrasts with the dark, artificial illumination inside the tunnel. If the camera's field of view includes both of these environments, then parts of the image will be useless—either saturated or too dark.

For law enforcement, cameras are usually placed at the front of a vehicle to record chase scenes. Because these images can be used in court as evidence, it is important to be able to see details such as a license plate or a driver's face, even if the camera is positioned toward vehicle headlights or toward the sun. Figure 1.4 shows typical luminance distributions.

In security applications, a situation similar to the tunnel scenario is encountered. A camera located inside a bank, for example, cannot simultaneously view the inside of the room and the outside of the building through the windows on sunny days. A security camera located at an intersection of a well-illuminated street and a shadowed street will only be able to view one of the two streets in a single image (see Fig. 1.5).

These situations can usually be solved by taking two images: one with the exposure adjusted for the bright part of the scene, and another with the exposure adjusted for the dark part of the scene. This approach makes sense in many situations but is not acceptable in others. If information is to be streamed to a control room, it will need to be split on two displays or somehow

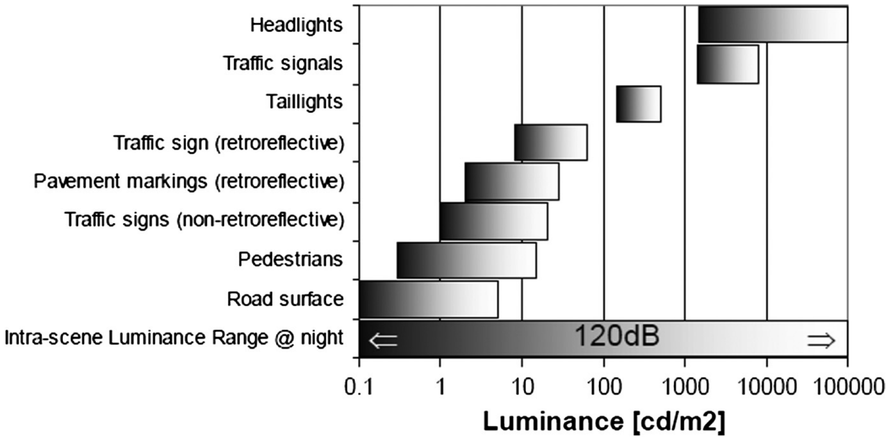

Figure 1.4 Typical luminance distribution in an automotive or law enforcement scene (reproduced from Ref. [67]).

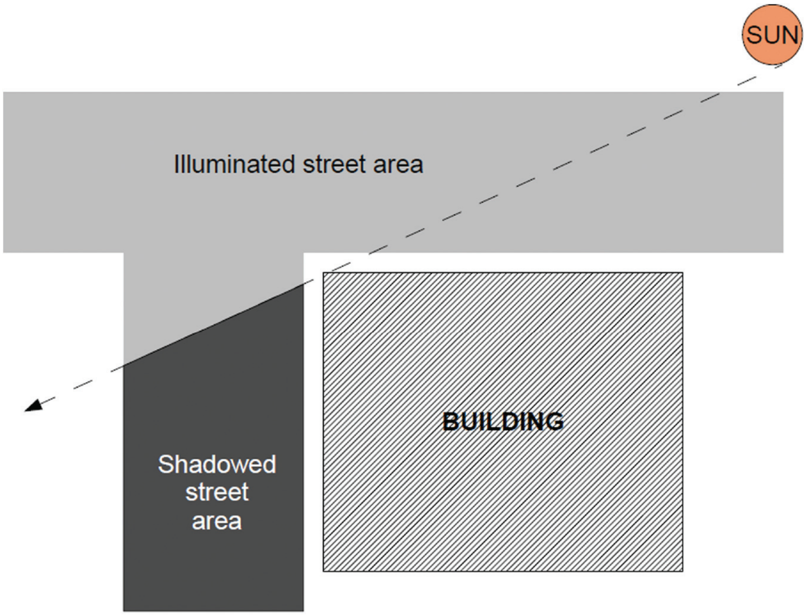

Figure 1.5 Shadowed-street situation.

combined. Splitting the data has a cost; combining it is an HDR technique described later. In other situations such as pedestrian detection, road sign detection, or lane departure warning, basing an algorithm on two images can be too slow for the application. The goals of these systems are meant to offer the highest possible degree of safety, yet at an affordable price. It may be too expensive to increase processing power and work with twice the frame rate. Reducing the overall frame rate affects the decision time, creating conflicts with the highest degree of safety requirement. Moreover, two different exposures can be insufficient for achieving acceptable results. In some situations, eight images or more are required to image the complete scene with enough detail.

Certain medical applications for diagnostic imaging require HDR imaging to see high-contrast images, which enable the detection of small variations in tissue coloration or local contrast. Another medical application that requires HDR is endoscopy because of the difficulties to provide uniform lighting.

Lighting issues are also present in projects related to oil- and gas-well inspections. Because of the nature of the environment and the narrow pipes, as well as the very limited embedded processing capability and communication bandwidth, multiple pictures of the exact scene cannot be made easily from the same vantage point, and so an automotive sensor with an HDR mode can be used.

Many night-vision goggles and other night-vision systems use amplification of light by photomultiplication. Similar results can be achieved by electronic processing. This principle works well in the dark, but saturation occurs as soon as a light source is present. It is possible to use a photon multiplier together with an HDR sensor to build a better electronic night-vision system.

HDR imaging is also important in certain defense applications such as target acquisition, dynamic target chasing, and missile interception. Future defense systems could involve lasers to blind an opponent's optical targeting or aiming systems. A very high dynamic range camera must be robust against such defensive equipment.

Barcode scanning sometimes suffers from the reflection of light caused by plastic covering the barcode. Although this can be solved by image processing for conventional barcodes, it becomes an issue of higher importance for quick-response (QR) codes.

The figures in this section show examples of luminance encountered in practical scenes. The ranges of luminance are often wider than the range that the film or sensor can image. A lux is the photopic unit of illuminance and emittance. It corresponds to the radiometric unit of W/m^2. The relationship between the lux and W/m^2 is the luminosity function, which is the standardized model of human visual brightness perception over wavelength. 1 lux $=$ 1 lm/m^2 $=$ 1 cd $\cdot$ sr/m^2. Candela per meter squared (cd/m^2) is the luminance photometric unit. A few numeric examples of luminance (from Ref. [76]) encountered in everyday situations are listed as follows:

- moonless sky: 3×10^{-5} cd/m^2
- full moon (0.2 lux): 0.06 cd/m^2
- low limit of human color vision: 1 cd/m^2
- living room (50 lux): 12 cd/m^2
- office lighting (500 lux): 125 cd/m^2
- cloudy sky (6000 lux): 1500 cd/m^2
- limit of human visual tolerance: 6000 cd/m^2
- sun: 2×10^9 cd/m^2.

Slide film usually covers a contrast of 8 f-stops or a contrast of around 250:1. Negative film is able to cover contrasts up to around 2000:1 or 11 f-stops. This means that negative film must be used to produce the best pictures of scenes in sunlight. Normal digital cameras range from below 8 f-stops to above 11 f-stops. Film is usually better than a medium-price-range digital camera.

Dynamic range is also an issue for non-imaging optical applications, such as the detection of reflected signals in depth measurement sensors, e.g., lidar and time of flight. Sun-load sensors must also deal with a wide input signal range.

1.2 High Dynamic Range Photography

HDR photography began in 1850 when Gustave Le Cray had the idea of combining a low-exposure negative and a high-exposure negative to reach a higher dynamic range. In 1945, HDR photography was required for taking pictures of the first nuclear explosions.

In the second and third quarter of the twentieth century, Ansel Adams (1902–1984) created very-high-contrast monochrome images using the zone system, a process he invented for selecting the proper exposure of glass plates (Adams was using large glass plates for their performance and resolution) and for selectively rendering the brightness of the final prints. Adams mostly photographed the American Southwest, more specifically, Yosemite National Park. His most famous photographs include *Monolith, the Face of Half Dome* (1927), *Moonrise* (1941), *Sierra Nevada* (1948), *El Capitan* (1968), *Yosemite and the Range of Light* (1979), and many others. The procedure used by Adams in his developing process, spatial manipulation, predated the multiple exposures and tone mapping used today in digital photography.

Modern commercial HDR imaging started in 1980 with the invention of the RGBE file format. At that time, HDR images were generated using local calculations. It was not until 1997 that images were generated using global calculation by Mann, Picard, and Debevec [76]. This is the approach used today, known as the radiance map or Debevec method (Fig. 1.6). This method produces very high dynamic range data that cannot be displayed, thus the tone-mapping theory compresses HDR data into a low dynamic range (LDR) dataset, which preserves the useful information (Fig. 1.7).

HDR photography is briefly discussed in this book at times because it yields great artistic pictures. However, the main purpose of this book is to

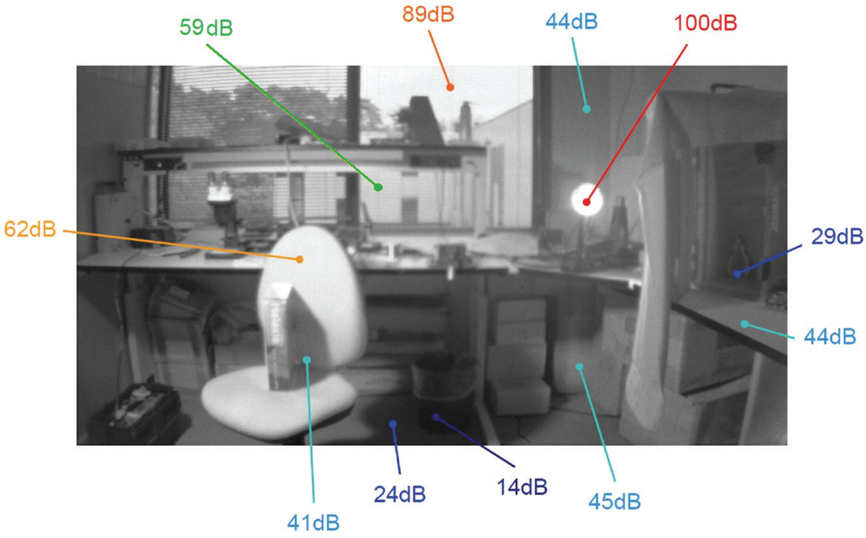

Figure 1.6 Image of an HDR scene and scene luminances relative to the sensor's lowest detectable level (courtesy of Melexis Microelectronic Systems).

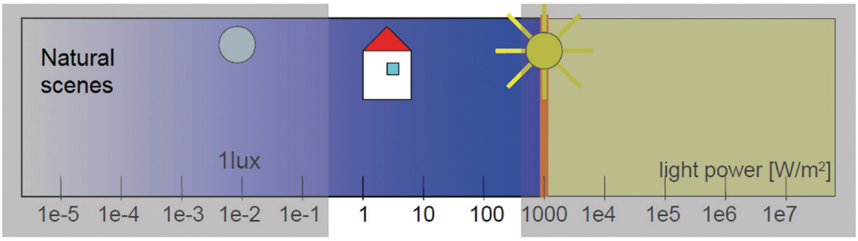

Figure 1.7 Comparison of luminance with the dynamic range of a common low dynamic range (LDR) sensor (courtesy of Dierickx [10]).

learn how to use HDR imaging in industrial machine vision applications or in similar applications that cannot be considered as photographic or artistic. This section shows a few high-quality HDR photographs (see Figs. 1.8–1.10 and 1.12); Fig. 1.10, for example, shows some halo effects that will be explained later in the tone mapping section. All images in the remainder of this book are related to industrial applications, or are used to show a specific effect.

To produce photographic HDR images of still scenes, a camera with a very fast exposure bracketing system (when the trigger is pushed, several images are acquired one after another with different exposure times) or a good tripod is required. See Fig. 1.11 for an example of a bracketing configuration and Fig. 1.12 for a tripod and remote-control shooting setup. A remote control is used to avoid any small changes to the position of the camera when the trigger is pushed or settings are changed. Many cameras today provide a USB or wireless connection to a laptop or smartphone. First, several images must be acquired of the same scene from the same

Figure 1.8 Village overlook at sunset, processed with open-source software and Fattal tone mapping.

Figure 1.9 Salisbury Cathedral, processed with Photomatrix Pro.

vantage point with several exposures. Then, the images need to be realigned and processed in specific software to extract an estimate of the radiance of the scene. Finally, the images are reduced to an 8-bit format using tone mapping.

1.3 Scientific Applications

Scientific applications often require images with extreme linearity and precisely known camera response to measure absolute and relative values of irradiance and luminance. As will be seen through the examples in this book, achieving this goal with an HDR imaging system is difficult, as the response curve of the sensor can be modified, and the algorithms applied can be

Figure 1.10 HDR photographs by Rachel Santellano (www.pixoto.com/rachels).

AEB .:1..0..1.:2.
White balance
Custom WB
WB SHIFT/BKT 0,0/±0
Color space sRGB
Picture Style Monochrome
Dust Delete Data

Figure 1.11 Exposure menu of the Canon EOS 40D, showing the exposure bracketing settings (AEB).

Figure 1.12 Equipment setup for HDR shooting.

nonlinear or might not preserve the sensor's responsivity. Optical veiling glare is also a significant limitation.

Recently, several CMOS image sensors featuring HDR modes have been specifically developed for scientific applications and tend to replace expensive, high-end CCDs.

1.4 High Dynamic Range, Wide Dynamic Range, and Extended Dynamic Range

Why are there different names? The term "high dynamic range" is generic; wide dynamic range (WDR) and extended dynamic range (XDR) are usually used by manufacturers when describing pixels or sensors. This book uses HDR in general to describe a method or a scene, and XDR to denote the extension of dynamic range compared to a non-HDR (also called low dynamic range) solution. To be complete, there is a last term, high dynamic range rendering (HDRR), used in the opposite process of displaying HDR images. HDRI stands for high dynamic range imaging.

1.5 Reducing the Exposure Time

When an image is overexposed, the exposure time can be reduced until the image looks less saturated. This can be done by changing some sensor registers or camera settings (see Fig. 1.13). In this figure, an image is previewed in the LCD screen of a camera together with its histograms, and the exposure time is adjusted until the histogram no longer saturates. With automatic exposure algorithms, this situation can be solved by using a negative-exposure compensation setting. An exposure compensation of -1 EV would mean a target brightness half that of the default mode.

Figure 1.14 shows the saturation of the image and the corresponding saturation peak of the histogram for the overexposed image. Figure 1.15 shows the clipping of the image in the darkest areas for the underexposed image. This underexposed image has no saturation. Figure 1.16 shows the

Figure 1.13 Adjustment of the exposure of a Canon EOS 40D based on histograms until the image is no longer saturated. The exposure time can be seen in the top-left corner of the screen.

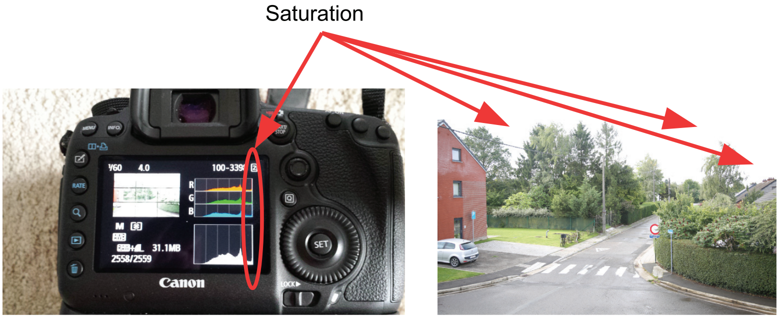

Figure 1.14 Saturation of the overexposed image.

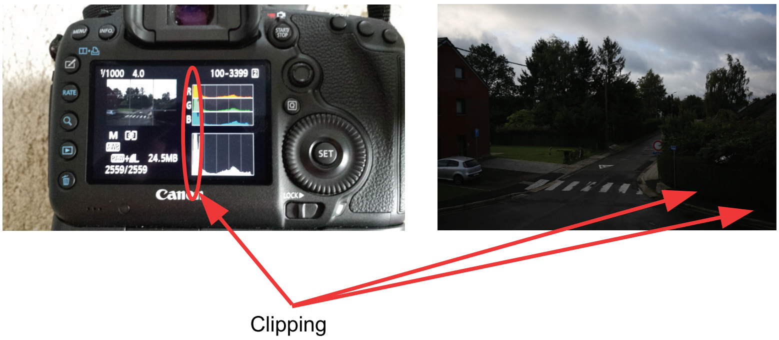

Figure 1.15 Clipping of the underexposed image.

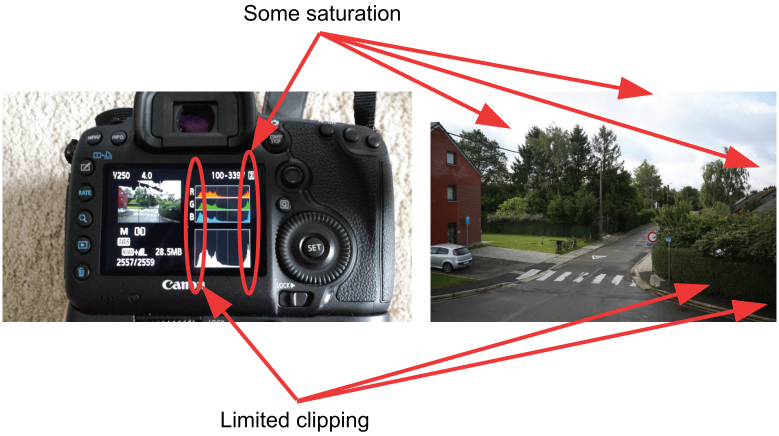

Figure 1.16 Best tradeoff image showing some saturation and limited clipping.

intermediate image that simultaneously has some saturation and limited clipping. It is the best possible tradeoff for this scene and this camera.

The image without histogram saturation clearly has more details in the highlights as the objects now appear, and it may seem that some software

Figure 1.17 Post-processing attempts of the last, non-saturated picture of Fig. 1.13 compared to an HDR version of the same scene taken with the same camera and the same saturation level. Top left: original image. Top right: gamma. Bottom left: strong sharpening. Bottom right: HDR merged and tone-mapped image.

post-processing would recover the dark details that have now disappeared due to the reduced exposure. However, this is not possible due to sensor and camera noise, as well as the analog-to-digital conversion. See Fig. 1.17, where different post-processing solutions are attempted but never produce the image quality of the corresponding HDR image.

1.6 HDR Applications That Do Not Require HDR Images

Sometimes the solution to an HDR problem can be found without using an HDR technique. For example, in lane-departure-warning applications, cameras have to detect road lane markings and estimate the path of the vehicle a few seconds in the future to determine if the vehicle will either stay in lane or cross a lane marking. Unwanted lane changes can happen due to driver drowsiness, and a warning signal (usually vibrations from the steering column) is provided to the driver. Because the road surfaces may reflect the sun and because the sun's location as well as other light sources are unknown, an HDR image sensor is usually used so that all details of the scene remain visible to the image processor, independent of the road condition or lighting situation. Lane markings are typically extracted from the HDR image; then a binary image is generated and processed to filter out everything but the lanes; and then calculations are made to generate a mathematical model of the lane curvature, the location of the lane center, and the vehicle's position in the lane.

This section will not describe the details of this complete process but rather show how the binary image of edges can be generated without using an HDR image, i.e., without merging exposures and without the use of an HDR pixel. Figure 1.18 illustrates the idea.

The proposed image sensor uses two different exposure times for the odd and even rows. A horizontal edge detection is performed, and the resulting data from the odd and even rows are merged, as if they were located at the same place. If the long-exposure row is saturated, no edges will be detected, but they will be detected in the short-exposure row. Conversely, if an edge is not detected in the short-exposure row because it is too dark and does not offer enough contrast, then the edge is likely to be detected in the high-exposure row.

During readout, the special proposed image sensor will perform the edge detection on the high- and low-exposure rows and will merge the results into a binary image (a single bit per pixel output that indicates whether or not an edge was detected in either of the two lines) by ORing the results of the

Figure 1.18 Road-marking-detection algorithm in an HDR scene without HDR imaging. Top left: even lines, long exposure. Top right: result of a horizontal edge detector on long-exposure lines. Middle left: odd lines, short exposure. Middle right: result of a horizontal edge detector or short-exposure lines. Bottom: combined and filtered binary edge image.

detection on the low-exposure and high-exposure rows. A more advanced approach uses 2D edge detection to reduce the edge noise.

But is this really not an HDR technique? Indeed, an HDR image is not reconstructed, but because either several images at several exposures or a specific sensor with dual exposure is required, it may still be considered HDR.

The same solution could be implemented purely in software out of two images at different exposures, but it would require memory and could suffer from significant motion artifacts due to the relative motion of the road between the two images. The two-image approach also requires a higher-speed sensor, as well as context switching.

There are other cases where an HDR image is not required, although the scene is HDR. For example, if the subject of interest is well exposed and the background saturates, the image will obviously not match the dynamic range of the scene, but this will not affect the result as the subject is well exposed, and no information is required from the background. In some applications, saturating the background can be a solution because it provides more contrast to the object to be inspected.

In many machine vision applications, the dynamic range requirements will be reduced by applying appropriate lighting or filtering techniques, such as spectral filtering, narrow-spectrum illumination, or polarizing filters. For example, barcodes are difficult to read under a plastic film because light strongly reflects on the film and the barcode vanishes into these reflections. The use of a specific light color and a specific type of light source solves the problem (see Fig. 1.19).

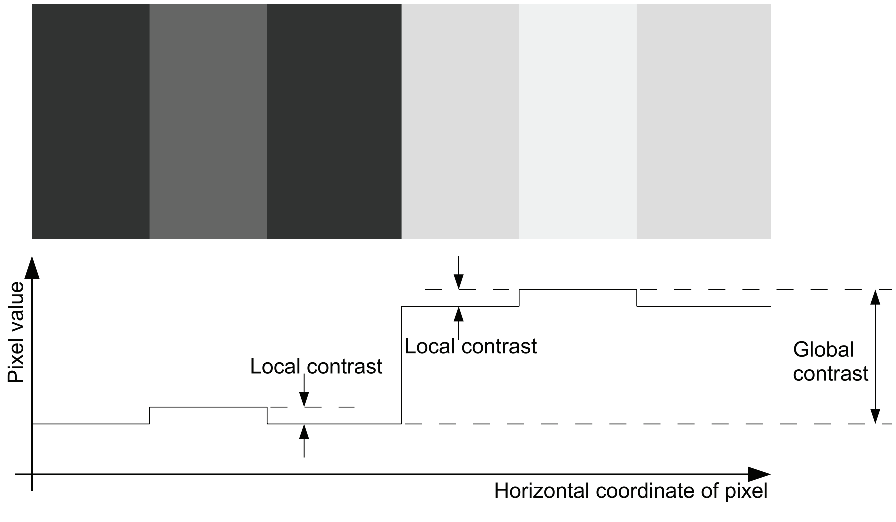

Figure 1.19 The scene's dynamic range of a barcode-reading application is reduced by the application of proper light sources and filters.

1.7 Image Histograms

A histogram is the representation of the distribution of numerical data. It shows the number of occurrences of each possible value of a statistical variable in a set of measured values. For images, the dataset is usually the pixel values of all of the pixels or a region of interest; it can be either the full histogram or only one histogram per color plane. Figure 1.20 shows an example of a histogram for a 5 × 5-pixel monochrome image coded on 3 bits (8 possible values from 0 to 7). Figure 1.21 shows the histogram of a real image. The total area of the histogram equals the number of pixels in the image.

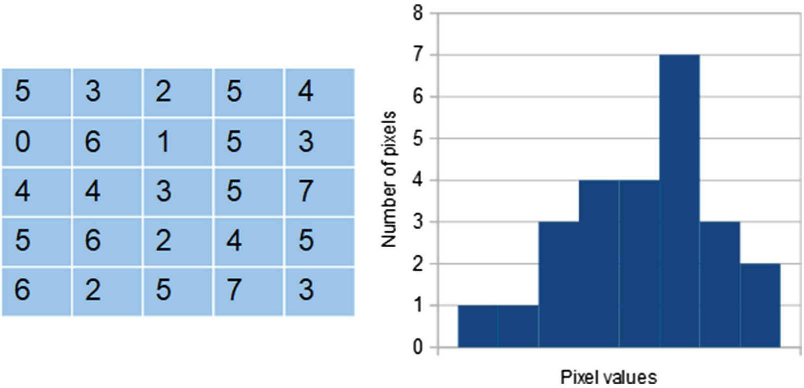

5	3	2	5	4
0	6	1	5	3
4	4	3	5	7
5	6	2	4	5
6	2	5	7	3

Figure 1.20 Sample 5 × 5-pixel monochrome image coded on 3 bits (left), and the corresponding histogram (right).

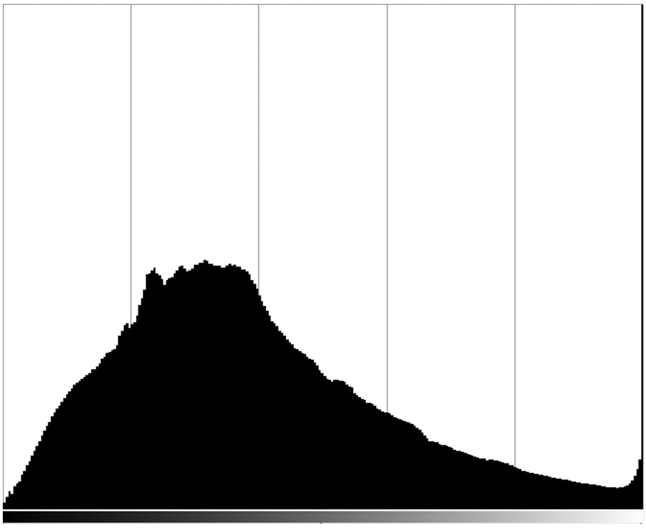

Figure 1.21 Histogram of a real monochrome image (generated with Gimp).

The histogram is very useful to determine if an image is overexposed or underexposed, as can be seen in Figs. 1.13–1.16. The histogram in Fig. 1.21 shows some saturation (the tall peak on the right), but the number of pixels in that peak is only a very small fraction of the total number of pixels, and therefore this saturation is acceptable. Histograms are thus very important in HDR discussions.

1.8 Outline and Goals

After this introduction, Chapter 2 begins with the definition and analysis of dynamic range from a mathematical point of view, and some important notions are introduced. To make this book understandable for most readers, some of the basics are re-explained when necessary.

In Chapter 3, possible system architectures are presented for HDR applications. This chapter primarily discusses the hardware methods used to extend the dynamic range of a camera system. Hardware methods are the methods used inside a pixel or sensor that do not fall into the software category. Any digital implementation of software belongs to the software category, although the software ends up on the chip. Hardware is covered in Chapter 3, and it corresponds to the HDR solutions for many industrial and military applications. The goal of Chapter 3 is to introduce many sensor solutions that exist or have existed commercially (and a few sensors that only existed in laboratories and/or publications) to familiarize the reader with XDR pixel design techniques, XDR performance and issues, and the mathematical techniques used to calculate the performance of such designs.

Chapter 4 introduces various software methods that can be implemented in a camera or on a computer. Due to the strong link between HDR imaging and tone mapping, tone mapping must be discussed, although it is not (strictly speaking) part of HDR extraction for machine vision applications. The world of HDR algorithms and tone mapping is very large, and Chapter 4 only gives a glimpse of it. The goal of this chapter is to familiarize the reader with basic concepts and some common algorithms.

Chapter 5 briefly discusses optical issues encountered in HDR imaging applications. Veiling glare and other optical effects are very important in photography and imaging, and the reader is encouraged to refer to specialized books and publications on this topic.

Chapter 6 reviews algorithms that can automatically control HDR exposure, as more complex sensors have more degrees of freedom and hence require more control logic. Once again, the concepts are only briefly introduced.

Chapter 7 discusses the file formats used to store images with XDR. The user interested in learning more details must refer to published standards.

Chapter 8 discusses the testing of HDR systems and sensors based on the International Organization for Standardization (ISO) and European Machine

Figure 1.22 HDR photgraph by Lofqvist [84].

Vision Association (EMVA) standards; the testing of HDR systems offers unique challenges that are addressed by these standards.

Chapter 9 draws some conclusions, presents a brief discussion of HDR figures of merit, and provides some final questions for students of HDR image sensing.

1.9 Defining a Camera

A camera is typically an imaging device made of an optical element, a sensing element, and some electronic circuitry. For this book, optical elements and electronic circuitry are assumed to be ideal, and the focus remains on the sensor and algorithms. A few very important optical and electronic limitations are introduced. In particular, it is assumed that there is no lens glare, although we will see later that this assumption is almost never correct and greatly affects dynamic range.