1

Introduction

The human visual system is a remarkable apparatus. It provides us with a three-dimensional perception of the world. Light is reflected from the objects around us. When the reflected light enters the eye, it is measured by the retinal cells. The information processing starts inside the eye. However, most information processing is done inside the brain. Currently, how this information is actually processed is still largely unknown. In computer science, we try to imitate many of the feats the human brain is capable of. Much work is still ahead if our algorithms are to match the capabilities of the human visual system. Among the many problems addressed by the computer vision community are structure from motion, object recognition, and visual navigation.

Color processing performs a very important role in computer vision. Many tasks become much simpler if the accurate color of objects is known. Accurate color measurement is required for color-based object recognition. Many objects can be distinguished on the basis of their color. Suppose that we have a yellow book and a red folder. We can distinguish the two easily because one is yellow and the other is red. But color can also be used in other areas of computer vision such as the computation of optical flow or depth from stereo based on color and shading. In this book, we will have an in-depth look at color perception and color processing.

1.1 What is Color Constancy?

Color is actually not an attribute that can be attached to the objects around us. It is basically a result of the processing done by the brain and the retina. The human visual system is able to determine the colors of objects irrespective of the illuminant. This ability is called *color constancy* (Zeki 1993). Mechanisms for color constancy also exist in other species, ranging from goldfish to honeybees (Tovée 1996). Color is an important biological signaling mechanism. Without color constancy, objects could no longer be reliably identified by their color. The visual system is somehow able to compute descriptors that stay constant even if the illuminant changes. Why is this remarkable? The receptors in the human eye

Color Constancy M. Ebner

only measure the amount of light reflected by an object. The light reflected by the object varies with the illuminant.

Let us follow the path of a ray of light to understand what actually happens. The ray of light leaves the light source and reaches the object at some particular point. The reflectance at that particular point, as a result of the material properties of the object, determines how much of the incident light is reflected. This reflected light then enters the eye, where it is measured. Three types of receptors exist inside the human retina, which measure the incident light. They absorb light with wavelengths in the long, middle, and short part of the spectrum. In order to accurately describe how much light is reflected by the object, we would have to measure the reflectance for the entire visible spectrum. The reflectance could be measured by illuminating the objects with white light or some other illuminant whose power distribution is known. However, a measuring device such as a camera does not know the power distribution of the illuminant. It does not measure the reflectance. Instead, it measures the product of the reflectance of the object and the amount of light hitting the object. Here we have two unknowns. We do not know the reflectance nor do we know the type of illuminant we have.

If we only take a measuring device with no additional processing, then the output will vary depending on the illuminant. Every amateur photographer has probably experienced this effect at one time or another. A photograph may look very different depending on the type of light source used. Sunlight or candle light produces much warmer colors compared to the colors obtained when a flash is used. It does not matter whether an analog or a digital camera is used to capture the image. Suppose that we have a yellowish light source. Light from the light source falls onto the objects of the scene. A white surface will reflect the incident light equally for all wavelengths. Thus, the white surface will appear to be yellow in the photograph. However, if we have a digitized image, we can process this image to obtain a better color reproduction. Many popular programs, such as Adobe's Photoshop, Google's Picasa, the GIMP, or Digital Arts Xe847 Photoshop plug-in, can be used to process images to obtain a color-corrected image.

In photography, this process is known as automatic white balance. Digital cameras usually have several options to handle white balance. If the type of illuminant that illuminates the scene is known, i.e. the illuminant is either a light bulb, a neon light, sun light, or a cloudy sky, then the white balance of the camera can be set to the appropriate option. The camera knows the color of these illuminants and is able to compute an image of the same scene as it would appear under a white illuminant. For some cameras, it is also possible to take an image of a white patch. The camera then uses the contents of this image to compute the color of the illuminant. Some cameras also have extra sensors to measure the ambient light in order to determine the color of the illuminant. The option automatic white balance is supposed to automatically select the best option in order to obtain an image as it would appear under a white illuminant. The ultimate goal is, of course, to just use the data that is available from the image, in order to obtain an image of a scene that looks exactly as it did to a human observer who took the photograph.

Human color perception correlates with integrated reflectance (McCann et al. 1976). Other experiments have shown that the human visual system does not actually estimate the reflectance of objects (Helson 1938). What is known about the visual system is that color processing is done in an area denoted as V4 (visual area no. 4). In V4, cells have been found that respond to different colors irrespective of the type of illuminant (Zeki

1993; Zeki and Marini 1998). Even though the visual system somehow computes a color constant descriptor, this color constant descriptor is not equal to the reflectance of objects. From a machine vision point of view, we are also interested in determining the reflectance of objects. If the reflectance was known, we could use the reflectance for color-based object recognition. The reflectance could also be used to segment scenes or for computation of optical flow based on color. In general, the reflectance information is very important for all aspects of color-based computer vision. For instance, an autonomous service robot should continue to work whether the environment of the robot is illuminated using artificial light or sunlight. If the computer vision algorithm uses the reflectance, which by definition is independent of the illuminant, then the environment of the robot (artificial versus natural lighting) would not make a difference.

Thus, there are two main roads to follow in developing color constancy algorithms. One goal is to determine the reflectance of objects. The second goal is to perform a color correction that closely mimics the performance of the visual system. The first goal is important from a machine vision point of view, whereas the second goal is very important for consumer photography. Since we are trying to determine the color of objects from just three measurements but do not know anything about the illuminant, or about object geometry, some assumptions have to be made in order to solve the problem. In the course of this book, we review many of the known algorithms that can be used to achieve better color reproduction. We also present new algorithms. All algorithms are evaluated in detail on different image sets in order to determine how accurate they are.

The problem of color constancy has fascinated scientists for a long time. Edwin H. Land, founder of the Polaroid Corporation, is one of the most famous researchers in the area of color constancy. In 1959, he performed a series of experiments with quite startling results (Land 1959a,b,c, 1964). Following this, he developed one of the first computational algorithms of color constancy (Land 1974, 1986a,b). Land had a major influence on the field of color constancy. His algorithms inspired many researchers, who then developed variants and improvements over the original algorithm.

1.2 Classic Experiments

Land (1962, 1964) realized very early on that the perceived color of an object depends on the rank order of the amount of light reflected for a given wavelength compared to the rank order of the amount of light reflected for another wavelength of the spectrum. It does not depend on the absolute values of reflected light. Land assumed that independent sets of receptors exist, i.e. one set for red, one for green, and one for blue, that operate as a unit to produce the perceived color. He named this system, which computes color constant descriptors, the “retinex.” The name retinex is a mixture of the words retina and cortex because at the time, Land did not know whether the color constant descriptors are computed inside the retina or whether the visual cortex of the brain is also involved. According to the retinex theory, the visual information processing starts with the receptors of the retina. Inside the retina, three types of sensors measure the light in the red, green, and blue parts of the spectrum. This visual information is then independently processed for the three color bands.

Land and McCann (1971) then developed a computational theory for color constancy, the retinex theory. In their experiments, they used a stimulus similar to the famous paintings

Figure 1.1 A Mondrian image similar to the one used by Land to develop his retinex theory. Land used colored sheets of paper and arranged them randomly. The resulting image reminded him of the abstract paintings that were drawn by the Dutch artist Piet Mondrian. This is why this stimulus is called a *Mondrian image*.

by Dutch artist Piet Mondrian. Later, Land (1983) noted that this stimulus actually looks more like a van Doesburg painting. Land and McCann used rectangular colored papers and arranged them randomly as shown in Figure 1.1. The colored papers were matted to reduce the influence of specular reflectance. Three projectors with sharp-cut bandpass filters were used to illuminate this Mondrian-like pattern. The first filter allowed only short wavelengths, the second allowed middle-length wavelengths, and the third allowed long wavelengths. The transmittance characteristics of the filters are shown in Figure 1.2. The amount of light emitted by each projector could be varied by a separate transformer.

All projectors were turned on and the transformers were set such that the papers of the Mondrian pattern appeared deeply colored. Also, the whites of the pattern had to be “good whites” (Land and McCann 1971). Next, a telescopic photometer was used to measure the light reflected from a particular area on the pattern. Land and McCann selected a white rectangle. They measured the reflected light, the luminance (see Table B.2 for a definition), using the telescopic photometer with only one projector turned on at a time. In other words, they measured the luminance reflected by the rectangle when the first projector was turned on. Next, they measured the luminance when only the second projector was turned on, and finally they measured the luminance with only the third projector turned on. In their experiment, Land and McCann obtained 60 short-wave units, 35 middle-wave units and 6 long-wave units for the white rectangle. Next, the telescopic photometer was turned to measure the luminance reflected from another rectangle. Land and McCann chose a rectangle that looked dark brown. The transformers were set such that the luminance reflected from the dark brown rectangle was equivalent to the luminance that was earlier measured for the white rectangle. In other words, the luminances were again 60, 35, and 6 for the short, middle, and long wavelengths, respectively. Even though the measured luminance was now

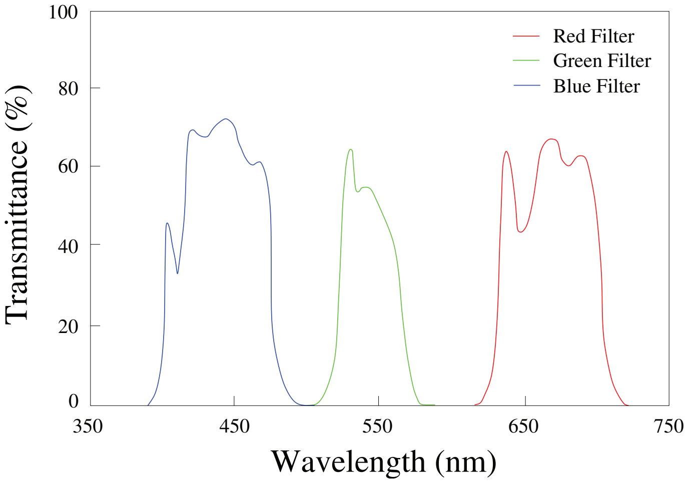

Figure 1.2 Spectral transmittances of the filters. The filters were placed in front of the projectors which were used to illuminate the Mondrian image. (Reproduced from Land EH and McCann JJ 1971 Lightness and retinex theory. Journal of the Optical Society of America 61(1), 1-11, by permission from The Optical Society of America.)

equivalent to that of the white rectangle, the rectangle was still perceived as dark brown. This experiment was repeated with different colored rectangles: bright yellow, blue, gray, lime, red, and green. For each rectangle, the transformers were set such that the measured luminances were 60, 35, and 6 for the short, middle, and long wavelengths, respectively.

In each case, the perceived color remained the same. In other words, to a human observer, the bright yellow patch remained bright yellow even though the measured luminance was equivalent to that of a white rectangle. Also, all the colors of the other rectangles of the Mondrian image also remained almost unchanged (in a few cases, there were some changes) when the transformers were adjusted to match the luminance from a specific rectangle to the luminance originally measured for the white rectangle. Land attributed the few cases where there were some changes to the shapes of the receptor responses of the retina. These receptors do not have the shape of a sharp-cut bandpass filter. Instead, they are Gaussian shaped with considerable overlap, especially between the red and the green receptors. Land later repeated the same experiment using a more standardized setting and obtained the same results (Land 1983).

With this experiment, Land and McCann vividly demonstrated that the perceived color of an object does not depend on the light reflected by the object. The perceived color depends on the reflectance, which specifies how much of the incident light is reflected. The reflectance determines the color of the object. The reflected light is essentially proportional to the product of the irradiance (see Table B.1 for a definition) and the reflectance of the object. A human observer is somehow able to derive the reflectances for the objects in view regardless of the illuminant used. In contrast, a digital or an analog camera can do no more than the telescopic photometer can. It only measures the reflected light. In order

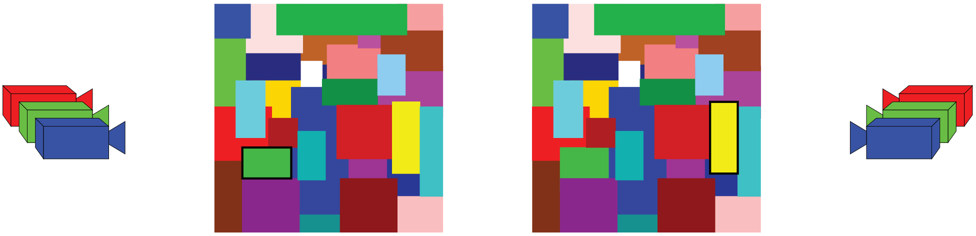

Figure 1.3 Experiment with two color Mondrians. Each Mondrian is illuminated by a set of three projectors.

to arrive at an image that is similar to the one observed by a human, we need to do some postprocessing of the measured data.

Land (1986a) and Land and McCann (1971) also experimented with small exposure time in order to see if the perceived image somehow depends on the amount of time the image is viewed. Observers also viewed the Mondrian through photographic shutters. This allowed Land and McCann to limit the exposure time to a fraction of a second. By limiting the length of time for which the Mondrian was seen, they were able to rule out that any adaptation mechanisms or eye motion could be involved in the mechanism for color constancy. The experimental setup was exactly the same as for the first experiment. The transformers were set such that the luminance reflected from a rectangle was equivalent to the luminance reflected from a white rectangle that was previously measured. The results were equivalent to the first experiment even if observers were only able to view the Mondrian image for one-hundredth of a second.

To see if the environment has an impact on the perceived color, Land and McCann also moved the rectangular patches around on the Mondrian. They report that the color sensation does not change significantly if the rectangle is moved to a new neighborhood where it is surrounded by different colored rectangles.

Land and McCann (1971) came to the conclusion that color perception involves structures of the retina as well as the visual cortex. On the basis of these experiments, Land and McCann developed a computational theory of color constancy, the retinex theory. Land later developed additional versions of his retinex theory. The retinex theory of color vision is discussed in detail in Chapter 7.

Land (1974) also performed an experiment using two color Mondrians. Each Mondrian was illuminated by three projectors (Figure 1.3): one projector with a short-wave filter, the second with a middle-wave filter, and the third with a long-wave filter. A white patch was selected from one of the Mondrians. The three projectors that illuminated this Mondrian were turned on and the amount of light reflected for each of the three color bands was measured. Next, a colored patch was selected from each Mondrian. The patch selected from the first Mondrian was green and the patch selected from the second Mondrian was yellow. The projectors were adjusted such that the amount of light given off by the two patches was equivalent to the amount of light measured for the white patch earlier. When all six projectors were turned on, observers reported that the color of the first patch was green and the color of the second patch was yellow. The reflected light was equivalent to the amount of light reflected from the white patch for both of them.

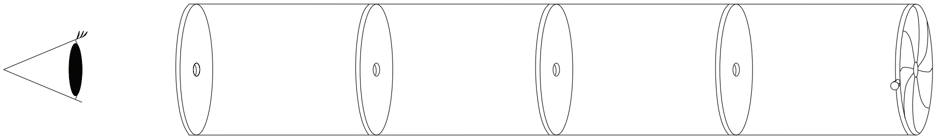

Figure 1.4 A long tube with internal baffles and an adjustable aperture. When observers view colored patches through such a tube, the patches look grayish-white. When the entire Mondrian is viewed, the patches appear to have a color that corresponds to the reflectance properties of the patch.

However, when both patches were viewed through viewing tubes as shown in Figure 1.4, the colored patches appeared to be grayish-white. The tubes were quite long with internal baffles and an adjustable aperture. The tubes were positioned such that the center of the green patch could be viewed with the left eye and the center of the yellow patch could be viewed with the right eye. The design of the tube ensured that only light from the center of the patches and none from the surrounding area entered the eye. When the whole Mondrian could be viewed, the patches appeared to be green and yellow. However, if the tubes were used to view the same two patches, they appeared to have the same color, i.e. grayish-white. From this, we see that the type of the environment of an object does have an influence on the perceived color of the object.

1.3 Overview

We now briefly describe how this book is organized. First, we have summarized Edwin H. Land's classic experiments on color perception. In Chapter 2, we have an in-depth look at the human visual system. Except for the wiring and response type of some cells, little is known about the actual algorithmic processing that is done by the human visual system. In Chapter 3, we have a look at the theory of color image formation. We will see how analog and digital color images are created. The course of light will be followed from the light source to the object and into the lens of the camera. At the end of the chapter, we arrive at an image that can be processed to achieve color constancy. Once an image is obtained, it also has to be reproduced, i.e. shown on a computer screen or printed on a piece of paper. Accurate color reproduction is discussed in Chapter 4. Color spaces are described in Chapter 5. Algorithms for color constancy are discussed in Chapter 6 and Chapter 7. Algorithms that assume a uniform illumination are covered in Chapter 6, whereas algorithms that do not assume a single uniform illuminant are covered in Chapter 7. Algorithms that assume a single uniform illuminant often try to estimate the illuminant. The estimate is then used to compute the reflectances for the object points. Such algorithms cannot be used for scenes with varying illumination. In practice, we usually have a scene with multiple light sources, for instance, a combination of natural and artificial light or simply several lamps placed at different locations of the room. Chapter 8 attempts to describe color constancy algorithms from a set of samples. Shadows are treated in Chapter 9, where we discuss methods for shadow removal and shadow attenuation. Chapter 10 explains how local space average color may be used to estimate the illuminant locally for each image pixel. In Chapter 11, we show how local space average color can be used to calculate a color-corrected image. Chapter 12 focuses on the computation of local space average color for an illuminant that varies nonlinearly over the image. In Chapter 13, the performance of all algorithms is evaluated on a set of images from a variety of sources. Chapter 14 compares the results obtained with the different algorithms to data obtained from experimental psychology. The most important points of this work are summarized in Chapter 15.