1

Introduction

1.1 WHAT IS COLOR GAMUT MAPPING?

On an average day most of us will come across a myriad of visual content generated using an ever-increasing variety of means. Already at the moment of waking up we may be presented with the flashing digits of an unwelcome alarm clock, accompanied by a sharp burst of light from a window or by a gradual emergence of shapes in the dark. During breakfast we may browse a newspaper or watch the news on a television. We may check our email using a personal computer or send a message, photo or video using our mobile phones and on the way out catch a glimpse of the latest artwork of our kids stuck to a fridge door or glance at some junk mail.

This scenario, which is so everyday as to be unremarkable, is an example of the ubiquity and variety of sources of visual content in our lives. We are regularly exposed to at least natural imagery, i.e. our homes themselves, whose reflective surfaces can be lit by natural or artificial light sources. Besides, we almost certainly view the output of traditional, analog imaging technologies such as drawings, print, photography and television and we regularly interact with digital imaging devices such as mobile phones and personal computers. In the process of work and leisure we also come across other imaging technologies, including digital projectors, cameras and printers. Many of us, therefore, are likely to have already used many, if not all, imaging technologies developed to date.

Reflecting on the diversity of visual content in our environment brings us to the observation that a variety of means can represent the same image and that these means, therefore, need to communicate among themselves. For example, we may wish to capture a moment from a birthday party (Figure 1.1) by taking a picture of it using a digital camera. Supposing we like the picture a lot we may also email it to our friends, send it to others via a mobile phone, print it out on a desktop printer, place it on a website, have a larger version of it printed in a copy-shop and include it in a presentation stored on a DVD and viewed on a television or projected onto a screen. Here, the same content (the scene from the birthday party) is present in at least 10 instances. Depending on how many friends we email and how many visitors come to the website, this number can be a lot larger.

Color Gamut Mapping Ján Morovič

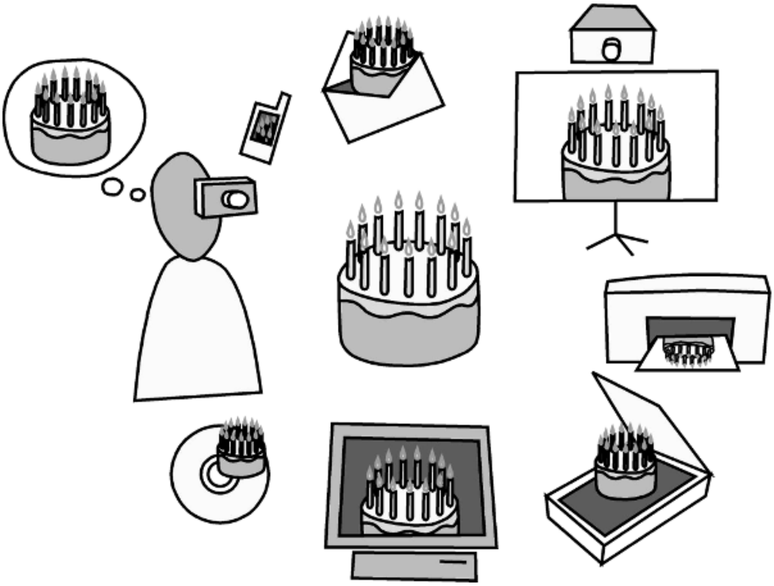

Figure 1.1 Birthday party images.

Next, let us think about how we would like the various instances of a given image to relate to each other. Clearly, the simplest and most immediate answer is that we want them all to look the same. In other words, we would like to see ‘the same’ when we look at the photo on the website and the print we made on a printer. Furthermore, we would like both of these to be ‘the same’ as when we looked at the actual scene during the birthday party. In some cases, though, we are less interested in an accurate record of an event and instead prefer to have an image that ‘looks nice.’ If the birthday party took place on an overcast day and everything looked a bit dull at the time, then we may still prefer for a more cheerful appearance to be represented by the images we took of the occasion.

Even though some differences between the various instances of visual content can be considered to be improvements, there are other differences that no one likes to see. Take, for example, the case of receiving the birthday party image and viewing it on your mobile phone in bright daylight. The image is likely to have much lower contrast than the original scene did, so much so that it could be difficult to see altogether. Viewing the image on your personal computer’s display could result in a darker result than how you remember the party. Printing the image might give a less colorful appearance than what was shown on your display. So, even if we want all instances of some visual content to ‘look the same’ or to have changes that make it ‘look nice,’ there are in practice many other differences between the various instances of an image that are undesirable.

Why is it then that different instances of a given image can look so different? The answer is a rather complex network of interactions among many individual factors, including the following:

1. *The instances of an image are viewed by different people.* As each one of us responds slightly differently to the light entering our eyes, there will be differences between the

experiences two people have when viewing a single image. Furthermore, as soon as we communicate about what we see, our experiences, skills and habits also play a role. When two people talk about a single image and even experience it in the same way, they are likely to express it differently.

2. *The instances of an image are viewed under different viewing conditions.* If we view two physically identical instances of an image in different environments, then they will look different. For example, viewing a television in the dark can give rise to a greater range of colors (i.e. greater contrast in images, more colorful parts of images, e.g. grass looking more vibrant) than when it is viewed in bright daylight, when it looks a lot duller overall. Note that in this case the television outputs the same image in both cases, but the dark versus bright environment changes its appearance significantly. This is yet more dramatic for images that reflect light (e.g. prints), where in the dark they too are dark and only when more light is present to view them do they acquire a clear appearance. In addition to how much light is present when viewing images, the background of the image is also important (e.g. what color the wall is behind a television), as is the distance at which it is viewed. Finally, it also matters what else is seen when instances of an image are viewed (e.g. whether the paper of a newspaper looks 'white' depends on whether we also see other kinds of paper at the same time) and how the instances are arranged (e.g. whether they immediately next to each other or far apart).
3. *The instances of an image are created using different technologies.* When the digital data from a camera are displayed on two different displays they are likely to look different, as displays can differ in terms of the materials they use for outputting color as well as in their settings. The digital data – which describe an image as a series of red, green and blue (RGB) values for a grid of spatial locations – give only relative instructions to imaging devices, i.e. the instruction may be to use 100% of a display's red and green colorants and 50% of its blue colorant. However, if the displays have different colorants, then following the same relative instructions will give different results. This constraint can, however, be overcome. The key to the solution is that it is possible to understand the relationship between digital data input to a display and the color appearance of the corresponding output (e.g. we can know what sending RGB = [100%, 100%, 50%] will look like on a given display and we can also work out what RGBs to send to the display if we want a certain color output from it). Then, to get two displays to look the same, we can take the RGBs we send to the first display, work out from them what the corresponding output colors will look like and from these color appearances work out what other RGBs to send to the second display. In other words, sending the same RGBs to two displays gives different colors, but sending different, appropriately chosen RGBs to each one can give the same color. Furthermore, the approach also extends to getting colors to look the same on different types of imaging device.

The above is a highly simplistic and incomplete sketch of why different instances of a given image can look different; more detail will be provided in Chapters 2 to 4.

One factor, however, was not made explicit in the above list. Namely, that, all else being equal (i.e. same viewer and viewing conditions), each imaging technology is capable of accessing only a specific part of all possible colors. For example, when viewing a television in a brightly lit room, it is not capable of providing a viewer with the experience of a very

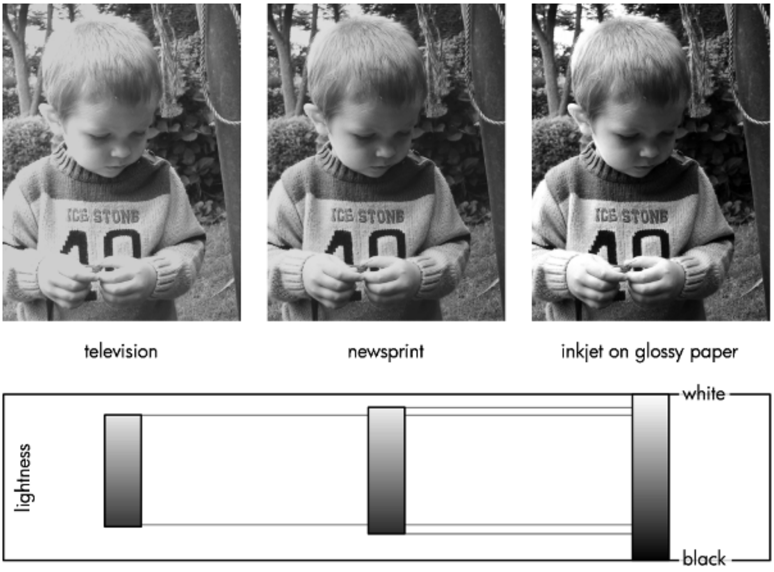

television newsprint inkjet on glossy paper

Figure 1.2 Getting dark shades in a brightly lit room. Top: a scene represented using different media; bottom: ranges of lightnesses possible on those media.

dark black. Looking at an image in a newspaper has such limitations too, whereas a glossy inkjet print does not and can give rise to the experience of dark blacks (Figure 1.2). A computer display is likely to be able to display more colorful, bright greens than are possible in print under the same viewing conditions, whereas print is more likely to be able to contain dark, colorful greenish-blues that cannot be obtained on a display.

As was proposed above, we tend either to look for colors being reproduced in the 'same' or a 'nice' way in all instances of an image. However, as individual imaging technologies provide different ranges of colors, our desire for sameness or niceness may in some cases be impossible. For example, start with an image that has the colors we want (e.g. a painting) and attempt to reproduce it by means that are unable to provide all of them (e.g. capture the painting using a digital camera and view the result on a television). While there will be parts of the painting whose colors we can replicate on a television, there are likely to be other parts that we cannot and decisions need to be made about how to deal with such colors.

Most important, these decisions can have a strong effect on the final image appearance and the same set of decisions is not likely to be appropriate in all cases. For example, when an image of a piece of clothing is to be reproduced in a printed catalogue, it is very important for that image to be as similar to the actual piece of clothing as possible. If there are some colors in the fabric that cannot be matched in the printed catalogue, then what matters is to represent them by colors that look most similar to the original. Alternatively, if a business presentation prepared on a computer display includes a pie-chart that has a segment in a bright, pure yellow, then it is preferred for such a yellow to be reproduced as bright and pure when printed, rather than as a yellow that is as similar to the original one as possible but which does not preserve the original's brightness and purity. Finally, when scanning and reproducing traditional holiday snapshots on a display, the color of the sky is typically preferred to be more vibrant than is possible in the photograph, rather than to be as close to it as possible.

What the above examples illustrate are both variety and, less directly, the potential for misapplication. Take, for example, the case of the printed clothes catalogue and consider the (mis)application of the solution from the holiday snap reproduction case. Clothes would appear more vibrant in the catalogue than they are in reality and this would be likely to result in customer disappointment, returned orders and reduced turnover. Conversely, if the catalogue solution were applied to holiday snaps, some displayed images would look dull and disappointing. Hence, careful decisions need to be made when dealing with the different color ranges available to different instances of an image.

More technically speaking: the color ranges mentioned above are called *color gamuts*, the changes made to images to deal with different instances of the image having access to different color gamuts are called *color gamut mapping* and the likelihood of gamut mapping being necessary in any given color communication is very high.

1.2 HISTORICAL CONTEXT OF GAMUT MAPPING

To counter the popular impression that gamut mapping is a modern phenomenon, it is worthwhile to consider its historical background before proceeding with a technical analysis. Such an insight into its origins will put gamut mapping into perspective and facilitate a better understanding of its expected effect. Since gamut mapping is the dealing with different color ranges being available to different instances of an image, its history coincides with that of man-made images. To create an image involves its reproduction either from another visible source or from the mind of its author and, hence, the addressing of differences between available color ranges.

As soon as humans created the first images, gamut mapping took place and it has been continuously practiced ever since (Figure 1.3). The earliest such images are Paleolithic cave paintings that date back 30 000 years and record hunting experiences on cave walls. The materials (i.e. colorants) used for creating these images, prime examples of which can be found in France (e.g. Chauvet-Pont-d'Arc, Lascaux) and Spain (e.g. Altamira), were charcoal and red and yellow ochres. When these images were created, there had to be at least an implicit consideration of how to represent the wide range of color experiences from a hunting event using the very limited range of available colors. Therefore, already these paintings are instances of the use of gamut mapping, albeit in a way that is implicit and certainly not like current approaches.

Following its Paleolithic beginnings, we can see a gradual but continuous expansion of the palette of color reproduction means, and a very brief look at their emergence will be taken next (Wikipedia, 2005a). Note that the development of the various means of creating color images often also provided access to greater color gamuts or to finer color variation within a given gamut. For example, the simple palette of the earliest cave paintings that consisted of three to five distinct colors became more extensive over time and was followed by the development of frescoes (i.e. 'painting in pigment in a water medium on wet or fresh lime mortar or plaster' (Wikipedia, 2005b)) in 15th century BC Greece. These frescoes were created using a much more varied palette of colorants and also allowed for

Figure 1.3 Examples of color reproduction from the Paleolithic age to the present. Digital original image reproduced by permission of Joe Kondrak.

an easier creation of transitions between them. Similar results were also achieved by painting onto dry walls, e.g. as seen in the tombs of Egyptian rulers from the 13th century BC onwards.

The 6th century BC saw the introduction of painted, enameled tiles in Babylon and the use of elaborate painted pottery in Greece. Mosaics and tapestries appeared around the 2nd century BC; and the earliest remaining carpet, which too is a means of reproducing color imagery, is considered to be the 5th century BC Middle Eastern *Pazyryk* rug excavated in Siberia. Frescoes from the 1st century AD can be seen to exhibit finer detail than earlier uses of the technique allowed, as do mosaics from the 5th and 6th centuries AD. In this period we also start seeing the first examples of illuminated manuscripts, which involve painting and the application of gold leaf onto parchment. The beginning of the second millennium AD heralds the use of elaborate, multicolored ceramics and stained glass windows, and imagery painted onto paper is preserved from subsequent centuries.

The earliest printed book that survives to this day is the *Diamond Sutra* from 9th century China and it is among the first examples of relatively large numbers of copies of visual content being produced. These beginnings of mass production are an important stage in the history of imaging, as they have two key implications. First, they lead to a process that is currently blossoming in the widespread availability and use of imaging technologies. Second, they introduce an additional source of imaging that is distinct from art and that has strong commercial and mass-consumption aspects. Had the creation of visual content remained exclusively in the hands of artists and craftsmen, the practice of gamut mapping might never have become explicit.

On the long road from the first use of book printing to the present and diverse variety of technologies, we also find the practice of painting onto silk, popular in 16th-century Japan, the invention of three- and four-color halftone printing by Le Blon in the 17th century (Pankow, 2005) and the widespread use of woodblock color printing in 19th-century Japan.

Around the beginning of the 20th century, color imaging experiences a rapid expansion thanks to the introduction of a variety of color photography techniques (starting with Lippmann's method, demonstrated in 1891, and followed by solutions that take us into the present day). The launch of color television in the late 1930s, of the first drum scanner in the 1950s, the first color computer display and inkjet printer in the 1970s and the first color digital cameras and color laser printers in the early 1990s continue the trend. The development of such imaging technologies, and especially the emergence of digital ones, is paramount to the explicit arrival of gamut mapping on the scene in 1978 (Buckley, 1978), as it is digital imaging that for the first time in history allowed for an explicit dealing with the differences of color ranges for the parts of an image, i.e. explicit gamut mapping.

The key point to take away from this overview is that gamut mapping has implicitly been present ever since man-made images were first painted on cave walls during the Paleolithic age and that its emergence in the 1970s is only the becoming-explicit of an ever-present process rather than its beginning. This realization is particularly helpful, in that it highlights the tight coupling between gamut mapping and the skills of artists and craftsmen who have practiced it for millennia, before its explicit form was taken over by mathematicians, scientists and engineers. An implication of this is also the need to confront scientifically developed solutions with a critique by those who create visual content professionally rather than only to focus on their scientific merits, which are in the foreground of contemporary gamut mapping research.

1.3 WHO IS THIS BOOK FOR?

The primary aim of this book is to help its reader to acquire the ability to choose from among existing, alternative ways of doing gamut mapping for an actual imaging task or product and then to implement them. More specifically, the focus will be on understanding and choosing existing gamut mapping solutions, as the other aspects of such an endeavor require familiarity with disciplines far beyond gamut mapping itself (e.g. color and imaging science, computer science, computational geometry, etc.). As regards these broader disciplines, a basic, and very much gamut-mapping-focused, overview will be provided of color and imaging science and of computational geometry. Experience with using at least some computer programming language and an introductory, university-level mathematics knowledge are highly recommended. Scientists, researchers, engineers, training instructors and students of undergraduate and graduate courses in color and imaging and related subjects are the primary audience for this book.

The following content will also address questions about gamut mapping that arise in the context of its use, rather than development and implementation. Hence, those who create color and imaging content (e.g. photographers, graphic designers, artists), those who provide technical services in the imaging industry (pre-press houses, print service providers, consultants), those who are involved in the preservation and distribution of color content, as well as undergraduate and graduate students of engineering and computer science courses will also benefit from the material presented here.

1.4 WHAT IS IN THE REST OF THE BOOK?

To aid both developers/implementers and users of color imaging technologies, Figure 1.4 provides a recommendation for how to use this book, depending on the reader's background.

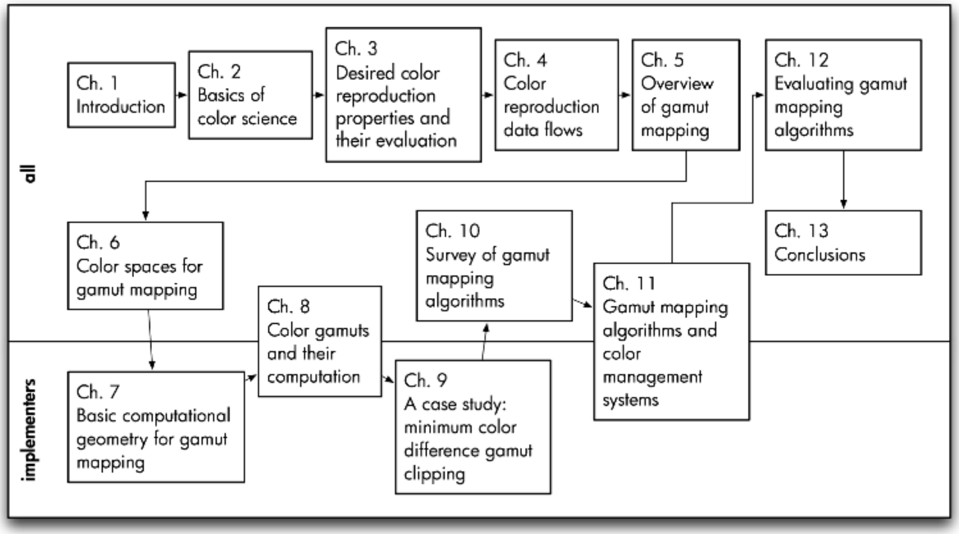

Figure 1.4 Book content overview.

Following this introduction, the basics of color science will be presented in Chapter 2, where the aim will be to give an overview of what needs to be taken into account in a system where gamut mapping is to be performed. This will include discussions of color appearance phenomena, viewing conditions, human vision, color-related physical properties, color measurement and quantification, color difference prediction and the basic principles of color imaging devices and media. The emphasis will be on providing a high-level outline and references to relevant literature. Finally, the chapter will conclude with applying the general overview to the question of what affects the perceived relationship of a pair of color stimuli (i.e. objects that evoke the experience of color perception).

Chapter 3 will take a look at various aims that color reproduction can have and at how reproductions can be assessed with reference to them. Psychophysical/psychometric versus measurement-based approaches will be introduced and an overview will be given of the *Commission Internationale de l'Eclairage's* (CIE's) *Guidelines for the Evaluation of Gamut Mapping Algorithms*. The concepts discussed in this chapter will be essential for a clear understanding of gamut mapping and will be referred to throughout the remainder of the book.

An 'under the hood' look at color reproduction workflows will follow in Chapter 4, where ways of representing original color information and the transformations it needs to undergo on the way from original to reproduction will be made explicit. Attention will also be paid to the importance of gamut mapping in different applications (e.g. what the role of gamut mapping is in newspaper printing versus press proofing on an inkjet printer).

Following the initial chapters that set the scene, Chapter 5 will take a detailed look at gamut mapping itself. The aims of gamut mapping, the types of gamut mapping algorithm (GMA – i.e. processes), their components and the factors affecting them will be introduced and different ways of looking at the role of gamut mapping will also be compared. Chapter 6 deals with the impact of the choice of color space in which gamut mapping is performed and Chapters 7 and 8 look at what color gamuts are and how they can be computed.

Armed with the content provided in the first seven chapters, a gamut mapping case study will be taken under the microscope in Chapter 9. For a specific color reproduction system, the journey of a color image will be traced from original to reproduction; also, how the image's colors change in the process due to the application of the minimum color difference gamut clipping algorithm will be looked at explicitly. This chapter will show the reader how the concepts introduced so far are applied to a concrete case and the effect they have will be scrutinized both in visual and underlying numerical terms.

Chapter 10 will then give a survey of gamut mapping work published to date, with an emphasis on considering the pros and cons of individual algorithms. Chapter 11 confronts gamut mapping research and development with its implementation in commercial color management systems, such as those based on the *International Color Consortium* (ICC) specification, the Windows Color System and sRGB workflows.

How GMAs can be evaluated, e.g. in terms of whether their performance differs for different kinds of images or in terms of how well they deal with smooth transitions in originals, will be covered in Chapter 12, and Chapter 13 considers current challenges and future trends in this area and the future role of gamut mapping in the wider context of imaging.

Finally, the book will not cover the following topics, apart from briefly mentioning their relationship to general gamut mapping:

- *First, the optimization of gamut mapping implementation.* As GMAs are defined in terms of the effect they are to have when applied to the colors of originals, there can often be considerable variation in the resources (e.g. processing, memory and time) that need to be used to achieve gamut mapping, depending on how it is implemented. Nonetheless, the focus here will be on looking at gamut mapping proper, in terms of the various ways in which sets of original colors can be adjusted to fit a reproduction color gamut, and the reader is advised to consult general texts on algorithm optimization where applicable.
- *Second, high dynamic range imaging (HDRI).* This refers to techniques that allow for the capture and subsequent processing of images containing large 'dynamic ranges,' i.e. ratios of the brightest and darkest parts of the image. Whereas, for example, computer displays have dynamic ranges of around 200:1, natural or computer-generated scenes can be of the order of magnitude of 100 000:1 (Hunt, 1995b: 787) and HDRI deals with such high dynamic ranges. HDRI addresses the capture, encoding and also the subsequent reproduction of high dynamic range (HDR) images and it is with regard to HDR image reproduction that there is a link to gamut mapping. However, as HDRI and gamut mapping currently deal with different domains (the former is applied to quantities of light in scenes/images and the latter to color appearance), there will be little discussion of HDRI in this book. Nonetheless, attention will be drawn to cases where there has been cross-fertilization between the two areas and also where there might be potential for it in the future. The reader interested in HDRI is advised to consult Reinhard *et al.* (2005).

Given the above sketch of this book's content, let us now proceed to discover in more detail the landscape at whose heart gamut mapping lives.