

CHAPTER 7

COLOUR MANAGEMENT

by Ján Morovič and Johan Lammens

Introduction

“The eye is the window of the human body through which it feels its way and enjoys the beauty of the world,” (Leonardo Da Vinci). And as that beauty is presented in full colour to the sense of vision, particular attention needs to be paid whenever colour is brought about, especially as colour can today be created by a plethora of means. Since these means encode, control and address colour in different ways, colour information needs to be managed.

The relatively recent emergence of new colour imaging technologies, e.g., digital cameras, printers, displays and projectors, their widespread availability and the use of colour information from one technology in another have made dealing with how colour is communicated a ubiquitous need. Take, for example, a family who use a digital camera to take pictures during their holidays, then view these pictures on their television and home PC’s display, print out some of the pictures on their desktop printer, share the pictures with their friends and relatives via a web site and get a commercial print service provider to produce a poster from one of their holiday snaps.

On closer inspection it can be seen that the above example involves: at least five distinct types of colour imaging technologies (each capable of addressing a different range of colour appearances), two or three operating systems, more than three ways of encoding colour information and at least six instances of interfacing colour information between technologies with their different capabilities and encodings.

The processes that translate and communicate colour information at such interfaces can be referred to as colour management. In other words, colour management is, for example, used to take colour information captured by the digital camera and translate it to make it suitable for display on a television or a computer display. Colour management is also used to take colour information viewed on a display and translate it for a desktop or commercial large format printer or translate it to make it suitable for sharing via the Internet.

In summary, colour management can be defined as “the process of providing a chosen relationship between colours generated using different imaging devices by translating colour information from a source device to colour information for a destination device.”

Colour reproduction objectives

A key question therefore is what the relationships between source and destination colours can and should be. While there is a natural inclination to say that "the colours should match," it will be shown next that this is rarely possible and even when it is possible it might not be desirable.

In his analysis of reproduction objectives, Hunt¹ presents a hierarchy of degrees of matching between a pair of colours (or colour images) and shows clearly why in most cases even a weak degree of matching is not possible.^a He starts with discussing differences in the *spectral* properties of source and destination colours^b and goes on to point out the effect of differences in light source intensity and chromaticity between the colour pair's viewing conditions, making a match in colorimetry impossible in most cases and undesirable in the rest. A match even in, e.g., relative CIE XYZ tristimulus values² between a source colour from a display with a white point of 9300K and a printed destination colour viewed under a D50 simulator would not preserve the source's appearance but result in a printed colour that looked bluish when compared with the displayed source.

The objectives that Hunt puts forward as practically meaningful are those of '*corresponding*' and '*preferred*' colour reproduction. Corresponding colour reproduction is equivalent to the CIE technical committee TC 8-03's 'subjective accuracy'¹⁹. In subjectively accurate colour reproduction the destination is "as close to the [source] as possible, this similarity is determined psychophysically and [there are] no image enhancing aims." In other words, this objective is about obtaining such a destination colour that looks as similar to the source colour as possible given the differences in the range of colour appearances (i.e. colour gamuts) of the source and the possible colour range of the destination under their respective viewing conditions. (See Chapter 11 CIE colour appearance models and CIE publication 159:2204³ for further details.) Note that this implies first, that a source image that looks unpleasant (e.g. has some defects) will result in a destination image that also looks unpleasant and second, that the destination device's colour capabilities might not be used to their full potential. An example of the second implication is the reproduction of a source image printed on uncoated paper by a destination image on coated, glossy paper that has a larger colour gamut. In this case the destination image would not make use of the entire available colour gamut if the 'subjectively accurate' reproduction objective were followed.

Since subjective accuracy is not what is needed in many colour reproduction scenarios, Hunt and the CIE also specify the '*preferred*' colour reproduction objective, which in Hunt's words intends to 'give a more pleasing result to the viewer' than the source did. For example, when printing holiday snaps there is less concern about the accuracy of representing captured images than about having 'nice' looking photos and the preferred colour reproduction objective is appropriate. Achieving this can include changes to the source that make memory colours more like their ideal prototypes⁴ (e.g. by changing the sky in an image's reproduction to

a Please, note that unless explicitly stated otherwise, colour reproduction will be used here to refer both to the reproduction of individual colours and of colour images.

b E.g. a CRT's red has the 'spiky' spectrum of emissions from a rare-earth phosphor while printers typically obtain red by combining their magenta and yellow inks, which are very much smoother spectrally.

INTRODUCTION

look ‘bluer’ than it was in the source), changes to the tonal and chromatic distribution of an image⁵ and other modifications of the destination image that adapt it to the destination gamut’s properties.

Viewing a pair of colours

Before addressing the question of how to manage colour information when interfacing it between two devices it is useful to look at a pair of colours generated using different imaging devices and consider the range of factors that contribute to their perceived relationship (e.g. whether they match or how they differ).




Figure 7 - 1: Viewing a pair of colours generated using different imaging devices.

The first thing to notice in Figure 7 - 1 is that there are many factors that are involved in determining the outcome of comparing the colour appearances of two stimuli and that if any of these factors change, the outcome of the comparison can change too.

COLOUR MANAGEMENT

More specifically it can be seen from Figure 7 - 1 that there are several factors besides the pair of colour stimuli that affect colour appearance:

1. the presence of other stimuli generated using the two devices or simply viewed alongside them (e.g.: Are the white points corresponding to the two devices visible? Is there a neutral more luminous than the two white points present when the colour stimuli are viewed? etc. – affects adopted white and hence perceived appearance attributes);
2. the background and surround against which the stimuli are viewed (affecting adaptation, involving simultaneous contrast^c and crispening^d effects);
3. the spectral power distributions, chromaticities, and illuminances of the light sources under which the stimuli are viewed (strongly affecting adaptation, setting limits to possible colours perceptible under them, determining the Hunt effect^e);
4. the distances from which they are seen (affecting perceived lightness and chroma^f);
5. the nature of separation between the two stimuli's viewing environments (i.e. stimuli can be seen simultaneously under mixed viewing conditions where the state of adaptation can be complex to determine, or at different times in which case the judgement about the relationship of stimulus appearances also involves memory^f);
6. the specific physiology of the individual observer (determining the particular colour matching functions, affecting perceptibility thresholds);
7. the experience the observer has with making the requested observations, comparisons and judgements (affecting perceptibility thresholds, judgement tolerances and repeatability).

The key points to take away from this analysis are: First, that the relationship between the ways in which two colour stimuli appear to an observer is not a property solely of those two stimuli but rather of a complex constellation of numerous states. Second, that it is not colour stimuli that are compared and judged, but the mental representations of their appearances, which clearly points to the inherently subjective nature of the colour comparison task.

These observations, however, lead to the question of how colour can be managed given that colour management only deals with interfacing and translating colour information between devices and therefore has potential to influence only a small part of a colour reproduction set-up – the colour stimuli. To address this question, at least the following four aspects need to

^c *Simultaneous contrast* refers to the phenomenon whereby the colour surrounding another affects it in the opposite direction of their difference (e.g. a darker surrounding colour will make the central colour appear lighter). See Chapters 11 and 12 for further detail.

^d *Crispening* refers to the phenomenon whereby a pair of similar colours looks more different from each other against a background similar to them than against a dissimilar background (e.g. the difference between a slightly lighter and a slightly darker than mid-grey colour will appear greater when they are seen against a mid-grey background than when they are seen against a white background). See Chapters 11 and 12 for further detail.

^e The *Hunt effect* describes the phenomenon of surface colours appearing to be more chromatic as illumination level increases. See Chapters 11 and 12 for further detail.

^f A time difference of even just 15 seconds introduces a ΔE^*_{ab} colour differences of around 5 units between the colour seen and the colour remembered and this difference slowly increases with time (de Fez *et al.*, 1998).

INTRODUCTION

be considered: First, that many of the factors have limited impact on the final judgement if their states change little (e.g. a viewing distance of 60 cm versus 70 cm will lead to similar experiences). Second, that if the comparison of colour stimuli is not simultaneous the limits of colour memory set wide thresholds on judgements and diminish the impact that factor changes and differences have. Third, that for colour critical tasks stimuli need to be compared under standardised or at least controlled viewing conditions (e.g. ISO 3664:2000 specifies how to view prints and transparencies in the graphic arts context). Fourth, that the complexity of colour reproduction set-ups introduces fundamental limits to what can be obtained.

In other words, colour can be managed closely only if other factors of a colour reproduction set-up, beyond the imaging devices, are also controlled. However, colour management can also lead to acceptable results even if such control is not possible either when comparisons are not done simultaneously or when the state of a colour reproduction set-up is close to what the colour management process assumes.

Conceptual stages of colour reproduction

Given the above discussion of colour reproduction objectives and the factors involved in the task of viewing and comparing colours, let us now turn to the conceptual stages that are involved in translating colour data from a source device into colour data for a destination device.

Before going into the details of the colour reproduction process, it is worth making a distinction between a *colour imaging device* and a *colour reproduction medium*. The difference between the two can readily be seen when considering a computer display device and a printing device. While in the former case the device is what is viewed to see the generated colour stimulus (i.e. we look at the stimulus generated by changing the properties of parts of the display), in the latter case it is a separate object that is altered with the help of the device and constitutes the colour stimulus that is viewed (i.e. the printed pattern on a substrate). Here the printer is an imaging device, the print is a colour reproduction medium and the display is both a device and a medium. An imaging device is involved in generating a colour reproduction medium, which is what is viewed in the end (e.g. the display, the print (but not the printer), the projection (but not the projector)).

Device colour spaces

Coming back to the process of translating colour information between a pair of devices (as it is these that can be controlled directly), the starting point is an encoding of source colours in terms of a device colour space of the source device (Figure 7 - 2).

COLOUR MANAGEMENT




Figure 7 - 2: Conceptual stages of colour reproduction.

The term *device colour space* refers to a space that does not in itself have a colorimetric interpretation but is simply used to address whatever colours can be generated using an imaging device. For example, the same red, green, and blue (RGB) values can be sent to a range of devices. Depending on which device they are sent to they will result in different colour stimuli. Hence such device RGB values do not in themselves represent specific colour stimuli but only allow for the addressing of the colour stimuli that can be generated on the device they are sent to.

A key advantage of device colour spaces is that they address all of the range of stimuli that a device can generate and nothing but that range and this makes them very well suited for addressing device colour capabilities. Of increasing relevance is also that device colour spaces provide a layer of abstraction between the inputs to a device and the colorants that the device uses. E.g., device RGB (dRGB) can be used to address all of the colour stimuli that can be generated using: a display with red, green and blue phosphors; a projector that in addition to RGB channels also has a white channel; a printer that uses cyan, magenta, yellow and black inks or also other inks like dilutions of some of CMYK. Each of these devices then deals internally with the assignment of colorant combinations to each device RGB triplet so that the device RGB space addresses all of its colour range. Analogously CMYK can also be used as a device colour space.

Device characterization and calibration

Given a set of device colour space coordinates (e.g. dRGB=[10,20,30]) from the source device, the first step of the colour reproduction process is to determine what stimulus the device will generate when receiving it as an input.

INTRODUCTION

This translation of device colour space inputs to stimulus outputs is the role of *device characterisation* and there exist a wide variety of models for this stage of the process that will be introduced in following sections of this chapter. In all cases, characterisation models have their parameters determined from pairs of (a) device colour space inputs and (b) colour measurements of stimuli generated by the device when receiving them. Given these parameters a characterisation model can then perform two predictions:

- for a device colour space input, predict the colour stimulus that would be measured if it were sent to the device (e.g. given $dRGB=[10,20,30]$ predict CIE XYZ tristimulus values that would be measured were it displayed on an LCD);
- for a stimulus that is to be generated using a device, predict what device colour space inputs to send to it (e.g. given CIE $XYZ = [30,20,10]$ predict what $dCMYK$ values to send to a printer to match it).

In addition to the level of accuracy achievable using a given characterisation model, a strong condition to its predictive powers is the state in which the device is to which it applies. If the device has not changed since measurements were made from which model parameters were determined, then all is well. However, if the device has changed – either because of changes to its components (e.g. ageing), to its environment (e.g. change of temperature, relative humidity), or to its settings (e.g. having a display's brightness setting increased) – then the relationships that the model attempts to represent may no longer hold and while its predictions might have worked in the original state, they now cease to hold.

To address the potential mismatch between the state in which a device has been characterised and the state in which it is at a later date, the process of *device calibration* is used. Simply put, the role of device calibration is to take a device in whatever state it is and restore it to a predefined state. Here the predefined state can be a standard one (e.g., $sRGB^8$ for displays) or the state of the given device at an earlier time. For a more detailed discussion of device calibration and characterisation see Bala.⁹

Colour appearance model

The next stage is to start with the prediction of the source stimulus and predict from it its colour appearance under the source's viewing conditions, as it is this appearance that is more closely involved in colour communication than the stimulus itself.

While there are several available models for predicting colour appearance, currently the most advanced model, able to predict the appearance of colour stimuli and suitable for use in colour reproduction, is CIECAM02 (see Chapter 11). In addition to CIECAM02, work on extending it to take into account some spatial phenomena and make it perform better for complex images (e.g. photographs, etc.) has also been done¹⁰ (see also Chapter 12) and is being further promoted within the CIE's TC8-08 on *Spatial Appearance Models*.

Finally, it is also important to be aware of the work done on dealing with the viewing of colour reproduction media where the state of adaptation is a mixture of component adaptation states. A scenario where this would be the case is the simultaneous viewing of a display and a print, where there is a significant difference between the white points of the two media (e.g. D93 for the display and D50 for the print). The colour appearance of the displayed colours will be different to what they would be if the display was viewed separately as a viewer would not be fully adapted to the display's white point and the actual, mixed state of adapta-

tion - affected by the simultaneous viewing of a print – needs to be estimated. To find more information about this issue, see the work of CIE TC 8-04¹¹ and for further detail on colour appearance models consult Fairchild.¹²

Colour and image enhancement

Given the appearance of the source colour under its viewing conditions, it is in some cases desirable to alter it, before considering its reproduction using a destination device. Note, that by definition this stage is not used if the reproduction objective is subjective accuracy as there it is indeed the appearance of the source colour that is the aim.

However, when preferred reproduction is wanted there can be reasons for changing colour appearance before its reproduction in the destination. One such change that can lead to more preferred reproductions is a change in hue whereby source hues are moved based on the primaries and secondaries^g of the source and destination media so as to make them more similar¹³. For example, if a source display medium's yellow secondary is greener than the yellow primary of a printed destination medium then keeping hue unchanged can result in a reproduction of a bright, pure source yellow as a darker, less chromatic greenish yellow. If, instead, the source hues are changes so as to move the source yellow towards the destination yellow then the reproduction can preserve more of the brightness and purity of the source colour at the cost of some hue change.

Finally, note that enhancements to the source content can also be performed in colour spaces other than those of colour appearance models. However, in any case their final output can be expressed in colour appearance terms and serve as the input to the next stage of the process.

Colour gamut mapping

Given the colour appearance that is desired in the destination – it can be either the appearance of the source or a modified, enhanced version of it – the next step is to ensure that it can be matched. To do this, it is first necessary to know at least the destination gamut and a means of determining it is required.

A number of techniques can be used to determine the colour gamut of a colour reproduction medium, including alpha shapes¹⁴, mountain range¹⁵ and segment maxima^{16,17}. In all these cases the starting point are colour appearances corresponding to a sampling of device colour inputs to the device that is involved in the colour reproduction medium whose gamut is to be determined. For example, when dRGB is used, a uniform sampling of all combinations of 10 steps per dimension (resulting in 10^3 samples) can be used. The gamut computation technique then uses this set of samples to generate a surface that delimits them from the rest of colour space.

^g A 'primary' of an imaging medium is a colour obtained by fully applying one of its colorants (e.g. a print's 100% yellow, a display's 100% red) and a secondary is a colour obtained by fully applying two of a medium's colorants (e.g., printing 100% of both yellow and magenta, displaying 100% of both red and green). The importance of these colours is that they play key roles in determining the shape of a medium's gamut.

INTRODUCTION

Next a transformation needs to be applied to all source colours that results in each of them ending up inside the destination gamut and this transformation is called *colour gamut mapping*.

Where the destination gamut is smaller than the source gamut, gamut reduction needs to be applied, which can either be a kind of clipping or compression. Clipping here refers to a gamut mapping where the source colours that are already inside the destination gamut are left unchanged and each of the source colours outside the destination gamut is mapped onto its surface (e.g. to the closest point on it). Compression on the other hand can change all colours – even ones that are already in gamut – so as to distribute gamut differences across a wider part of colour space and allow for preserving more of the colour differences that were present in the source than is possible with clipping. Finally, where the destination gamut is larger than the source, gamut expansion can be applied to make use of some of the additional colour space.

Note that all three kinds of gamut mapping may be applied to a single source–destination gamut pair as the destination may be smaller in some parts of colour space (e.g. around red), but larger in other parts (e.g. around cyan). For a more detailed look at gamut mapping see Morović.¹⁸ For the evaluation of gamut mapping algorithms see CIE¹⁹.

Completing the process

Starting from source device colour values, the process has taken us to a colour appearance that is desired in the destination medium and it is next necessary to determine what stimulus, under the destination’s viewing conditions, has that appearance. A colour appearance model is used in the inverse direction^h and results in the stimulus that is to be produced in the destination medium. The characterisation model of the device that generates it is used again in the inverse direction to predict the appropriate device colour inputs, which are then sent to the destination device. The pair of colours – i.e. the source, which is the starting point of the reproduction process, and the destination, which is its end – are finally viewed and their relationship is judged with respect to the chosen reproduction objective.

The ICC colour management framework

While the previous section outlined the conceptual stages of a colour reproduction process, we will now consider how such a process can be implemented in practice. This in turn leads directly to the *International Color Consortium* – ICC, whose colour management framework is currently the *de facto* standard, at least as far as the reproduction of still images is concerned.

The ICC was established in 1993 by eight imaging companies, “for the purpose of creating, promoting and encouraging the standardization and evolution of an open, vendor-neutral, cross-platform colour management architecture and components.”²⁰ To this end, the solution proposed by the ICC is one where the colour reproduction process is divided into two transformations: First, a forward one that takes device colour data and transforms it into a colori-

^h I.e. the forward direction is to predict appearance from information about a stimulus and its viewing conditions and the inverse it to predict a stimulus given a desired colour appearance and viewing conditions under which it is desired.

COLOUR MANAGEMENT

metric description for specific viewing conditions (called the *Profile Connection Space – PCS*). Second, an inverse one that takes such a colorimetric description and transforms it back into device colour space data.

Colour interchange between devices is then achieved by being able to perform both parts of the transformation for each of the colour reproduction media among which colour is to be managed (Figure 7 - 3). The parameters based on which the forward and inverse transformations are performed for a given colour reproduction mediumⁱ are stored in a data file referred to as the ‘ICC device profile’ and its detailed specification as well as the specification of the entire architecture can be found in ICC²¹. An overview of what the key parameters are for different types of imaging devices will be discussed in following sections of this chapter.




Figure 7 - 3: Overview of ICC colour management architecture.

The PCS, through which all colour communication takes place, is defined by the ICC as “the reference colour space in which colours are encoded in order to provide an interface for connecting source and destination transforms” (see pp. 8 of ICC²¹). The colour spaces that can be used for the PCS are CIE XYZ and CIE LAB for a reference viewing environment, defined for the graphic arts by ISO 3664²² viewing condition P2 standard (D50 light source; 500 lux illuminance; 20% surround reflectance).

As colour reproduction objectives are essential to colour reproduction, the ICC too specifies alternatives for them and refers to them as *rendering intents*. The four rendering intents defined in the current version of the specification are:

First, the *media-relative colorimetric* intent re-scales in-gamut tristimulus values to map a medium’s white point to the PCS white point and is useful for reproductions between media

ⁱ Note that a specific profile is needed for each colour reproduction medium rather than just for each imaging device. E.g., a printer printing on plain paper will need a different profile to that same printer printing on glossy paper.

INTRODUCTION

to which observers are fully adapted. A popular use of this rendering intent is also in conjunction with black point compensation – BPC²³ where the source luminance range is linearly scaled to the destination luminance range before gamut clipping is performed.

Second, the *ICC – absolute colorimetric* intent leaves tristimulus values of in-gamut colours unchanged and is useful for reproducing, e.g., spot colours and for proofing.

Third, the *perceptual* intent is useful for a preferred or pleasing reproduction of images, particularly pictorial or photographic-type images – especially where source and destination media are substantially different. To allow for more control in providing preferred colour reproduction, the ICC specifies a reference medium for this rendering intent. This medium is an ideal reflection print with a specific dynamic range and its purpose is to allow for improved results when performing gamut mapping. There are also proposals within the ICC to define a gamut for this reference medium to provide further control over the re-rendering and gamut mapping process that has to occur in two stages via the PCS.²⁴

Fourth, the *saturation* intent is also vendor specific, involves compromises such as trading off preservation of hue in order to preserve the vividness of pure colours and is useful for images containing objects like charts or diagrams.

In summary, colour transformations in the ICC framework are performed between devices on the basis of device profiles via the Profile Connection Space and rendering intent choices and when colour data is to be communicated it is necessary to provide them alongside the colour data itself (Figure 7 - 4). For further detail see the ICC web site (<http://www.color.org>), which also includes useful white papers on a number of colour management topics.




Figure 7 - 4: ICC workflow sketch (RI – rendering intent).

COLOUR MANAGEMENT

sRGB colour management

An important, complementary approach to the ICC colour management architecture, where the colour imaging behaviour of each device is characterised with reference to colorimetry, is to base all colour communication on a single device-related, but colorimetrically defined, colour encoding.

Specifically, colour communication workflows can also provide good results by taking two decisions: first, that RGB will be used to communicate colour information between devices and second, that RGB will be given a unique colorimetric interpretation – e.g. sRGB.⁸ Each device then does the best it can to either encode its native colour information in sRGB so that the result is pleasing (e.g. scanners, digital cameras) or to provide pleasing colour output given sRGB input (e.g. printers, displays, projectors).

The key properties of this approach are that only RGB content gets passed between devices and that each device internally does the best it can to relate the colours it generates or captures to sRGB. A clear advantage of such a set-up is that it is very simple and transparent to other elements in colour reproduction workflows, such as operating systems and software applications (Figure 7 - 5). The flip side though is that only a single reproduction objective can be followed by each device, which is significantly more challenging than if a specific reproduction objective were communicated alongside colour data.




Figure 7 - 5: sRGB workflow sketch.

Nonetheless, the approach works very well where pleasant colour reproduction is the aim and where users do not need to or want to customize the reproduction process. Finally, it is also worth bearing in mind that sRGB workflows can also interface with ICC ones via an sRGB ICC profile and that other colorimetrically defined encodings can be used instead of sRGB to set up analogous workflows (e.g. Adobe RGB²⁵).

Challenges of colour management

Even though the above introduction to colour management has already presented a high degree of complexity, there are a number of additional challenges that make control over the relationships of source and destination colours even more challenging.

INTRODUCTION

First, the relationship between device colour space data and corresponding stimuli for a given device is not constant. Instead, as components of devices vary temporally and also spatially across the device's imaging area and as components of devices are replaced from time to time, so does the way their device-specific colour data relates to stimulus colorimetry. Hence characterisation models and profiles are valid only for the state in which the device was when it was characterised or profiled. E.g. a profile generated for a device today may not describe it well in two weeks' time.

Second, the output of a given device may be viewed under multiple viewing conditions. E.g. a print may be viewed in a graphic designer's studio under daylight, in a press-room under a D50 simulator and in an end user's home under tungsten light. However, that print is made with a single specific viewing condition set-up in mind and its appearance in other conditions is not considered.

Third, differences between the visual systems and, more importantly, the colour preferences and colour judging experiences of individual observers: As soon as colour is to be generated for viewing by more than a single person, it becomes unlikely that each viewer will interpret it in the same way and control over the effect of the generated colour becomes limited.

Fourth, increasing variety of reproduction technologies and their gamuts. As colour image content is being re-used between devices as dissimilar as newsprint, cell phones, inkjet prints on glossy media, and laser projectors, the magnitude of differences that colour management needs to bridge increases and some of its components are stretched very far.

Fifth, proliferation of colour encodings and colour management implementations. Hand in hand with the increasing variety of reproduction technologies comes also a greater variety in the means of encoding and managing colour and it becomes a significant challenge to ensure that the way a colour reproduction set-up is arranged is consistent and communication is effective. E.g. as colour management itself is possible at different points in a workflow (e.g. in a software application, an operating system, an imaging device), it is easy for it to happen more than once and therefore for transformations to be applied to other data than is appropriate. sRGB can be wrongly assumed to be device RGB, device CMYK can be thought to be SWOP CMYK, etc.

Does colour need to be managed?

Merriam – Webster's 11th Collegiate Dictionary list a number of potentially relevant definitions of "manage": to handle or direct with a degree of skill; to make and keep compliant; to treat with care; to work upon or try to alter for a purpose²⁶. All of these would seem to imply active involvement on the part of the manager, and a certain degree of difficulty. Indeed, looking at Figures 7 - 2 and 7 - 3 one gets the impression that this is difficult business, while the set-up of Figure 7 - 5 looks a lot simpler. One often hears and reads that colour management is hopelessly complicated and something for eggheads, whereas a simple system like the one based on sRGB has proven to work well, and does not need colour management. As will be clear from the introduction, we do consider both ICC-based and sRGB-based systems to be examples of colour management, albeit different in their apparent complexity and degree of user involvement needed. We will illustrate this below. One might also think that

COLOUR MANAGEMENT

colour management is a monster created by the move from analogue to digital systems on one hand, and by the simultaneous move from closed and single-vendor to open and multi-vendor systems on the other. We will argue against this view of the genesis of colour management.

Despite their apparent differences, the ICC and sRGB systems as illustrated in Figures 7 - 3 and 7 - 5 actually have two important things in common: first, the definition of a common interchange “colour language”, and second, the translation of that common language into or out of each individual device’s private “colour dialect”.^j Some of the differences reside in the presence or absence of:

- explicit colour “dictionaries” (profiles) or descriptions of a device’s “private colour language”;
- an explicit “translator” component (colour management module) to convert one device’s colour language into another’s, and
- the use of explicit parameters (rendering intent) to modify that translation for some particular purpose.

As explained in the introduction, the common colour language in the ICC architecture is either CIE XYZ or CIELAB, which in the case of the *perceptual rendering intent* is output referred²⁷ to an *ideal reflection print* as seen under well-defined viewing conditions. The common colour language for the sRGB architecture is also output referred, but in this case to an *ideal CRT monitor* as seen under well-defined viewing conditions. The colour space of that ideal monitor is itself described in terms of the CIE XYZ coordinates of its primaries, among other things. So even though ICC can be thought of as more print-centric (or traditional Graphic Arts centric) and sRGB can be thought of as more monitor-centric (or perhaps computer centric), they both use an explicit output-referred exchange language which in both cases relies on fundamental CIE colour spaces for its definition.

In both architectures, participating system components (devices or software applications) need to be able to translate their own colour languages into or out of the common one (the PCS). This requires device calibration and characterization, as explained in the introduction. In the ICC architecture, the result of characterization is stored in an explicit and standard format, which can be exchanged freely within ICC compliant systems. A distinction is made among input, display, and output devices, with corresponding differences in the type and complexity of characterization formats (profiles) used. Input profiles are unidirectional (device to PCS only), and can use simple or more complex device models (see below). Display profiles typically use rather simple device models that can be inverted, and can hence be considered bi-directional (device to PCS as well as PCS to device). Output profiles tend to use more complex device models and are required to be bi-directional, among other things to enable proofing (the simulation of one device on another, different one). In the sRGB architecture, device characterizations (or the corresponding colour transformations) are typically fixed and built directly into the devices themselves, which makes them invisible from a user point of view – yet they do exist. In some cases, particularly CRT monitors, no explicit de-

^j Wittgenstein might not take lightly to the suggestion that devices speak a private language or dialect, but we will ignore this for the sake of the discussion.

vice characterization or colour transformations are needed since the device's colour space and behaviour correspond closely to those of the ideal monitor that the common colour language is referred to. This has obvious efficiency benefits for devices like monitors with their high data throughput needs. The built-in characterizations or colour transformations of the sRGB architecture are typically unidirectional only, which follows from the fact that they are built-in, hidden, and not exchangeable.

The “translator component” is known as CMM or Colour Management Module^k in the ICC framework. Its task is to connect two (or more) colour profiles via the defined connection space (converting between CIE XYZ and CIE LAB as needed), accept rendering intent parameters specifying which of the defined ICC rendering intents to use, and construct a “colour world” transform that can subsequently be used to transform colour coordinates from input to output colour spaces. Actual transformation of colour coordinates typically involves interpolation using different size and precision lookup tables, and possibly different interpolation algorithms (tetrahedral and tri-linear being the most popular ones). The CMM is a “pluggable” component in the ICC architecture, i.e. in principle it can be replaced at will by any other ICC-compliant CMM. In typical sRGB implementations (there is no standard architecture defined for sRGB systems), the “translator component” is a fixed part of a device or software program, and invisible to both the system and the user. Yet it performs much the same task of transforming between input and output colour spaces and interpolating values using lookup tables and interpolation algorithms.

As mentioned, ICC systems allow (and in fact require) explicit parameters and components to perform colour management operations. The profiles that describe device colour behaviour, the rendering intent parameters needed to perform conversions and the CMM itself are all explicitly defined using standard formats and interfaces, and meant to be replaceable and exchangeable. sRGB systems, on the other hand, are more “closed” systems in the sense that only the interchange colour space is explicitly defined, standardized, and accessible, but little else. Nevertheless they perform very similar tasks.

Analogue Colour Management

How did things work in the analogue age, and did colour management exist, or was it needed at all? Let's consider only a few representative examples: traditional silver halide photography, and traditional graphic arts production of printed publications. Traditional silver halide based photography involves (very schematically) the following steps:

- An image is captured by exposure to light of a film substrate coated with a photosensitive emulsion. Light interacts with the photosensitive chemicals in the emulsion, altering their properties in function of the amount of light received (exposure), and the spectral characteristics of the light as filtered through coupled colour filters. Both negative (print) and positive (transparency) films exist and have been widely used.

^k Sometimes also glossed as “Color Management Method”.

COLOUR MANAGEMENT

- The film is “developed” using chemical means, which “fixes” the changes brought about by exposure to light, and makes the film insensitive to further changes, i.e. makes it no longer be light sensitive. This results in a negative or positive daylight viewable representation of the image on the developed film.
- The developed film is used to create an image on paper by shining a light through it and projecting the resulting image onto photosensitive paper. The process is very similar to the one that created the negative or positive: the paper contains photosensitive chemicals coupled with colour filters that form a latent image after exposure, which is then developed and fixed through chemical means, resulting in a daylight viewable image on paper (and the paper no longer being photosensitive). Positive and negative photo papers exist, analogous to positive and negative films.

Is there any colour management in all this? There is indeed, although not in the form of digital colour profiles and CMMs. Starting with the capture process, one has to choose an appropriate film type for the lighting conditions (daylight or indoors), which mostly refers to the spectral composition of the illuminant. Each type (and often brand) of film has to be developed with a specific process. The paper type used must match the film characteristics. Photosensitive chemicals and coupled colour filters must be carefully designed and manufactured to exacting standards. Device calibration is ever present, in the form of exposure measurement and control for the camera and enlarger; temperature, timing, and concentration control for development agents; and so on. The main reasons that the process seems simple and automatic (at least to consumers) are that it has been under development for about two centuries now, and a large degree of standardization has taken place. Digital colour management is still in diapers by comparison, and standardization is still in its early stages. Nevertheless, an all-sRGB digital photo capture, editing, and printing system as is common for consumer systems today comes quite close to the ease of use and thoughtlessness of the famous Kodak “Brownie” systems. In Brownie days, all the complexity and science was hidden in film R&D and manufacturing plants on one end, and in photo processing laboratories on the other end – but there can be no doubt that it was there. If sRGB systems can be compared to Brownie consumer systems, could ICC systems perhaps be compared to the much more sophisticated and difficult to use professional analogue photography systems? Perhaps so, but there is no reason to assume that ICC systems could not be made as easy to use as sRGB systems, if that were the objective. The move from analogue to digital photography has opened up many new possibilities, many of which still have to be worked out and “brought under control” (standardized).

Now let’s have a brief look at traditional (analogue) graphic arts print production systems, using the watercolour reproduction scenario described in the next section. Completely photographic systems would follow much the same flow as described above, so let’s concentrate on electronic but analogue graphic arts systems. The basic steps involved would be (very schematically):

- The original artwork is scanned in a high end (drum) scanner, which directly produces analogue CMYK separation signals (voltages), as calculated by an embedded analogue “colour computer”.

ANALOGUE COLOUR MANAGEMENT

- The separation signals are used one at a time to expose pieces of photo sensitive separation film, one for each ink colour to be used in printing. Both positive and negative separation films have been used, analogous to traditional photography.
- The separation films are developed and fixed, and used to expose photosensitive printing plates, much like photographic negatives or positives are used to expose photosensitive paper.
- After development and fixing, the printing plates are mounted on a printing press, inked, and used to sequentially print each ink colour onto the final substrate (paper or other).

Is there any colour management going on here? There certainly is. The colour computer that calculates separations is converting images from RGB to CMYK colour spaces, which is one of the things that CMMs do in the digital domain. It typically does not use explicit device profiles for this, although some models have used programmable lookup tables, not unlike the lookup tables used in ICC profiles today. Everything related to film and plate exposure and development uses the same kind of chemistry-based colour management that we hinted at above. Process control (calibration) is essential to the proper functioning of a printing press, as is precise control of the composition of printing inks, papers, and other consumables used. The process may seem less complex and more robust than modern day heterogeneous digital systems, but the main reason for this is that they were essentially closed systems, relying on a large degree of standardization. Again the move to digital systems has opened up many more possibilities, many of which still have to be worked out and “brought under control” (standardized). But we see no reason why digital colour management ought to be more complex or less robust than its analogue predecessors, given the necessary time and effort.

Watercolour reproduction scenario

To make the following, detailed discussion of colour management more concrete, we will introduce an example scenario, which will serve as a backdrop for the remainder of this chapter. This scenario (Figure 7 - 6) revolves around the production of promotional material for an exhibition of watercolours and involves the following stages: scanning of a watercolour original; viewing and editing of the scanned image; page layout of a poster and a leaflet for the exhibition and their proofing and production.

Please, note that the focus in the following sections will be on the colour management aspects of the scenario rather than on a comprehensive description of the technologies and processes that it would involve. Hence when addressing the scanning of the watercolour original, the discussion will revolve around how to relate scanned data to the original artwork rather than how to best scan it (i.e. scanner requirements, settings, original handling, etc.).

COLOUR MANAGEMENT




Figure 7 - 6: A watercolour reproduction scenario's workflow.

Original to scan

The first task in our watercolour scenario is to obtain colour data that encodes the appearance of the original watercolour artwork and to colour-manage this process. A scanner, which illuminates a transmissive or reflective original and then samples and digitises the light modulated by it both spatially and spectrally, is therefore used.

To relate data obtained from a scanner to the original's colour properties, it is necessary to characterise or profile the scanner. This requires that it be in a state in which its output is stable and it is also advisable to calibrate it. The outcome of having a stable and characterized scanner is that an original's colour properties will be predictable from scanned data, which is needed when the original is to be reproduced using an output imaging device.

In general it is advisable to first warm up a flat-bed scanner before use as the built-in light source's output varies most in the initial minutes of operation and its output subsequently stabilises. When accuracy is important, it is also worth determining the uniformity of scanner response across its scanning area, as some scanners tend to have non-uniformity, parallel to the scanning sensor, in the region where scanning starts¹. As characterisation models and ICC profiles assume that there is a fixed relationship between original colour and scanned data

¹ A cause of this can be the fact that some scanners only have their light source on when scanning and they therefore switch it on shortly before each scan. The first part of the scanned image can therefore exhibit warm-up changes.

irrespective of scanning time and spatial position it is important to know how a given scanner departs from that assumption and to optimise its use relative to the accuracy requirements of a given task.

While calibration is, strictly speaking, not necessary (i.e. the scanner could be characterised in an arbitrary state), it is highly advisable as it allows for the scanner to be returned to the calibrated state at a later time and the characterisation model (or ICC profile) set up initially can be re-used. If calibration is not done, characterisation needs to be repeated frequently and as calibration tends to involve fewer resources, it is advisable to do it instead. Here it can also be beneficial to include linearization as part of the calibration process – i.e., to define the calibrated state to involve a linear relationship between scanned data and some property of the original (e.g. luminance or lightness, depending on how the scanner will later be characterised).

Characterisation models in general require pairs of colour stimuli and corresponding device data, and this is also necessary for scanner characterisation. Taking a set of uniform colour patches, such as the patches of a photographic IT8.7/1 chart²⁸, and both scanning (and averaging) them and measuring their CIE XYZ tristimulus values provides such pairs, from which model parameters can be determined. With its parameters determined, the characterisation can then predict what colour stimulus was presented to the scanner, given RGB data obtained from it.

Challenges of scanner characterisation

Characterising a scanner, alas, presents a number of fundamental challenges that complicate the simple picture given so far.

First, there are a number of issues that arise from the fact that a scanner records only three values for each spatial location in the source, while the colour related properties of the source (i.e. its spectral reflectance or transmittance) are not three-dimensional. This in turn means that a whole range of source spectral properties will result in the same scanned RGB values (Figure 7 - 7). In other words, there is a many-to-one relationship between source spectra and scanned RGB values and the question of what the source was like, given particular scanned values has many answers.

This many-to-one nature of the scanning process has three key implications: (a) That comprehensive characterisation models need to predict a source colour set given a single RGB response – such a set is referred to as a *metamer set*.²⁹ (b) That if the spectral response of the scanner is different from that of the human visual system^m (as represented, e.g., by the $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ colour matching functions (see Chapter 3) then metameric problems arise from: the scanner-human observer difference directlyⁿ and from differences between the light source used in the scanner and the light source under which the source is viewed.

^m Or, more precisely, if the scanner responsivities are not a linear combination of the human visual system's sensitivities.

ⁿ As a result of such observer metamerism, the scanner may record different RGBs for parts of the source that look the same to a human observer and the same RGBs for parts that look different.

COLOUR MANAGEMENT




Figure 7 - 7: Sample spectral reflectances resulting in identical RGB responses on a given scanner under D50 (PM Morovič, pers. comm. 5 March 2004).

Where there is a difference in scanner and human spectral responses and where the characterisation model is not of the many-to-one kind (e.g. ICC profiles), there is also the problem of different characterisation parameters being needed for sources of different kinds (i.e. different types of photographic paper, digital output, original artwork materials, etc., all need different parameters). In other words, a characterisation model that is accurate for some photographic material is likely to be inaccurate for a watercolour original.

Second, characterising scanners also presents a problem specific to input imaging devices that follows from the difficulty of fully sampling the inputs to the device. In general it is important to sample all possible inputs to a device and then determine characterisation model parameters from these samples and from corresponding outputs of the device.³⁰ While this is not difficult with most output imaging devices, where the inputs are digital values that can be generated easily, with input imaging devices, the range of possible inputs is the range of all possible colour stimuli. Even restricting ourselves to reflective and transmissive materials, it becomes unfeasible to have a set of samples that cover the gamut of all possible colour properties (i.e. the Object Colour Solid³¹). As a consequence a characterisation model will only be valid and have known properties over the colour gamut of the samples that it is based on and its predictions for stimuli outside that gamut will be unreliable. This is particularly a problem where there is a need to have a characterisation model that works for a wide range of source kinds and less of an issue where samples are possible across the full gamut of a particular kind of source (e.g. when only photographs made on specific photographic paper are scanned and a sampling of that paper's gamut is available).

Third, a number of scanner properties (e.g. tone response, quantization, noise) result in the scanner having an effective gamut³² that is more limited than the Object Colour Solid (OCS). This means that for certain sub-volumes of the OCS the scanner is not able to differentiate between distinct colours and records them as being the same. The most direct example

of this is the effect of the scanner's dynamic range, which results in variation below a given luminance not being recorded and being mapped onto a single value instead.

Scanner characterisation models

Scanner characterisation models can be categorised in terms of two criteria: model-based versus empirical⁹ and many-to-one versus one-to-one. Here model-based characterisation, which tends to be many-to-one, attempts two parallel stages of the scanning process in its computational structure, whereas empirical approaches treat the process as a black box and are concerned primarily with prediction accuracy.

Model-based approaches use information about the spectral responsivities (or sensitivities) of a scanner, the scanning light source and the scanner's tone reproduction^o to relate RGBs to scanned spectra. The forward direction (i.e. from surface reflectance to colorimetry) is in this case significantly simpler than the inverse as it involves a reduction in dimensionality from n to 3, where n is typically at least 31, and there are several detailed descriptions of how to perform the mapping (see pp. 315–316 in Bala⁹).

The inverse direction of a many-to-one model is, however, more complex as it needs to take a single RGB value and predict from it the set of all possible spectra that could have resulted in it under the given sensor and light source conditions. One such model is *Metamer Constrained Color Correction*,³³ which predicts a convex metamer set for a given set of scanned values and also provides methods for choosing a single representative of the set where necessary. This is important in an imaging context as, even though the RGB^p to spectral reflectance (and therefore also XYZ) relationship is a one-to-many one, in the end a given scanned RGB needs to be represented by a single XYZ, which is then further transformed in a colour reproduction system. Furthermore, having the entire set of possible XYZ that correspond to a given RGB and knowing the XYZs obtainable on another device (e.g. a display) also allows for the choice to be made in a way that optimises cross-device reproduction. In the absence of such information a choice can be made on statistical grounds (i.e. as some XYZs occur more frequently than others in nature or a particular original medium).

Empirical models, on the other hand, simply require a list of source XYZs and scanned RGBs based on which their parameters are computed. The simplest model, which is an application of Yule's Masking Equations³⁴ to scanners, first involves linearization of the scanned data by applying 1D transformations to the scanned RGBs so as to make them linear in terms of XYZ. Next, a 3x3 matrix is computed (e.g. using a linear least-squares approach) and applied to the linearised RGBs to predict XYZs from them. This model works well only if the scanner sensor responsivities are close to being linear combinations of the $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ colour matching functions (as the model assumes). As this assumption often does not hold, an extension of the simple model is to use higher-order matrices that allow for non-linear relationships between RGB and XYZ³⁵. Alternatively neural networks can also be used to em-

^o What is meant by tone reproduction is the relationship between device colour space values (i.e. RGB here) and corresponding colour or density attributes. E.g. the relationship of scanned RGB values of a reflective greyscale's steps and the greyscale's lightness values.

^p In this section the X , Y , Z and R , G , B tristimulus values will be written in a short form as XYZ and RGB.

pirically model scanners^{36,37} as can direct interpolation from the data obtained from a characterisation chart^{38,39}. Finally, note that many of the empirical techniques used for scanner characterisation can (by their general nature) be used for modelling other types of imaging devices as well.

Scanner ICC profiles

The ICC defines two types of profiles for colour input devices: *Three-component matrix-based*, which allow for the storing of the simplest model's parameters introduced in the previous section (i.e. three 1D look-up tables for linearization and a 3x3 matrix), and *N-component LUT-based* input profiles (see pp. 22 of ICC²¹). The latter type's most important component is an ICC *AToB0Tag* data structure, which stores parameters for the following sequence of transformations: input data \rightarrow 3x3 matrix \rightarrow 1D input LUTs \rightarrow multi-dimensional LUT \rightarrow 1D output LUT \rightarrow output data. In the case of scanner profiles the input data is scanned RGB and the output data is XYZ or LAB for the Profile Connection Space (PCS), the multi-dimensional LUT is three-dimensional and the rendering intent is perceptual. Analogously *AToB1Tags* contain parameters for the colorimetric and *AToB2Tags* for the saturation rendering intents.

More specifically, the LUT contains PCS values for an even sampling of the input RGB space and a key challenge here is to populate all entries of the LUT as it is often the case that the RGB – XYZ pairs obtained by scanning a characterisation target do not cover the entire RGB cube and some form of extrapolation is needed. Furthermore the CIE XYZs or LABs stored for each LUT entry are not simply measurements of source colours (or values interpolated from them) but values that represent a chosen rendering intent.

On a related subject, it is worth noting some issues when using ICC input profiles with images from digital cameras. Most cameras (or camera raw processing applications) colour render images to some standard output-referred colour encoding, like sRGB⁸ or Adobe RGB²⁵, so the profile attached should be the appropriate colour space profile. It may also be possible to obtain “raw” camera RGB, which has not undergone any colour space transformations. In this case, it will be possible to use a new kind of ICC input profile in the future that will interpret the camera RGB values. Such profiles, however, will be image, not camera specific. The colorimetric intents will depend on the scene illumination, and the perceptual intents also on the desired colour rendering of the scene. While it is mechanically possible to use the colorimetric intent of a raw camera RGB profile to convert to a standard scene-referred colour image encoding, such as RIMM RGB⁴⁰ or scRGB,⁴¹ doing so today may result in interoperability issues and that is something that the ICC is currently addressing. Note that if the colorimetric intent of such a camera RGB profile is used to convert to an output-referred colour encoding, this effectively indicates that the result is the desired colorimetry on the encoding reference medium. However, it is also possible to consider the initial result of such a transformation as the starting point for manually applied colour rendering.

The key point to take away from this section is that ICC profiles are simply parameter containers and that the challenge is primarily in the computation of these parameters, which can in turn be done by performing various stages of the colour reproduction process (e.g. characterisation, colour appearance modelling, colour enhancement, gamut mapping). The LUT for the perceptual rendering intent will therefore contain values that are a re-interpretation (or re-rendering) of the PCS colorimetry that corresponds to given scanned

RGBs and that give a perceptually more pleasing reproduction on other media (e.g. changes to contrast, saturation, etc.). On the other hand, the colorimetric rendering intent's LUT will represent the colorimetry of a scanned original and will only deal with presenting the scanned data under PCS conditions as opposed to the viewing conditions in the scanner.

Scanned watercolour

In terms of our scenario, the outcome of its first stage is an RGB image obtained by scanning the original watercolour, and to allow for the controlled reproduction of the scanned data, also an ICC profile of the scanner.

The key challenge here is to generate a profile that gives accurate predictions of the scanned original as it would ideally have to be based on watercolour colour patches painted using the same paints and paper as the original that is to be scanned (and also aged in the same way as the original has aged). However, as this is likely not to be possible, an attempt would have to be made to at least use a characterisation target that is as similar to the original watercolour as is practical. The scanned data together with the scanner's profile provide a colorimetrically based description of the original that will be the input to following stages of colour manipulation and reproduction.

Scan to display

The next stage of the watercolour scenario is to view the scanned image on a display to then be able to edit it and incorporate it into the page layout of a poster and a flyer. As the scanning stage of the scenario results in RGB data representing the original watercolour and as the associated scanner profile also provides a colorimetric representation of the scanned data (either with a perceptual or colorimetric rendering intent), the next colour management task is to provide an appropriate re-rendering of the scanned image on a display.

What will therefore be needed again is a characterization of the display that relates displayed colorimetry to the digital input that brought it about as this will be the basis for determining what display inputs to use to represent given scanned pixels.

Just as was the case with the scanner, a display too needs to be warmed up, as its output is different soon after being switched on from what it is like an hour later. Unlike with scanners though, the issue of sampling inputs to a display is trivial as there are no obstacles to driving it with digital values that sample the full range of possible digital inputs to it.

Challenges of display characterisation

There are two key challenges to successfully characterising a display. First, that displays suffer from non-uniformity, either spatially across the displayed area (on a CRT the luminance of output from the centre of the display can be up to 50% greater than from the edges for the same digital input) or as a function of viewing angle (colour output of LCDs can vary dramatically with viewing angle changes).

Second, as displays are typically not viewed in complete darkness, it is also necessary to take into account the impact of ambient illumination on the appearance of display output and this is a much more serious challenge than the first. The difficulty of dealing with the impact

of ambient conditions is twofold: First, there is a challenge to knowing how to take into account both the white point of the display itself and the white point of the ambient illumination to arrive at an effective white point that controls the colour appearance of display stimuli. Here there has been significant progress made by CIE TC8 - 04, who recently published a report on how to deal with such mixed-adaptation conditions¹¹. Second, however, there is a practical challenge of knowing what the ambient illumination is like and also of a colour reproduction set-up staying up to date with ambient illumination changes (e.g. in a studio where daylight is used, that daylight is likely to change) and this is much more difficult to address. In practice displays are typically profiled in complete darkness and if these displays are viewed under other conditions then the mismatch between profiling and actual conditions contributes to shortcomings in cross-media colour reproduction.

Display characterisation models and their implementation in profiles

As displays are in colorimetric terms simple devices whose channels' output is additive (i.e. the separate CIE XYZ values of the red channel's output and the green channel's output can simply be added to get the CIE XYZs of their combined use), their characterisation models too are simple and can be very accurate. To predict the colorimetry of a display's output given the digital input to it, display characterisation models perform two stages of transformation. First, the digital values of each channel are transformed to be linear in terms of luminance and this can be done using a range of models that are specific to different display technologies. E.g. the GOG model^{42,43} is a popular solution for CRT displays, the sigmoidal model is suitable for LCDs,⁴⁴ while the PLCC model⁴⁵ works for any display technology as it simply used 1D LUTs for the task. The second step in all cases is to transform the luminance-linear RGBs to CIE XYZ and this can be achieved using a 3x3 matrix obtained either directly from the RGB primaries of the system or optimised across the whole gamut. All of the above models can also be inverted to predict RGB display inputs from desired CIE XYZs. As more than three primary systems are emerging there is also a growing literature proposing solutions for their characterisation (e.g. Murakami *et al.*⁴⁶).

Specifically for the characterisation of CRTs there is also a CIE recommendation⁴⁷ that provides full details of how to apply Berns' GOG model and also specifies how to take into account flare and other factors affecting display output.

Turning to the implementation of display characterisation models in ICC profiles we find that they can be of the same two types as scanner profiles. Here the *Three-component matrix-based* type is most common and it allows for the storage 1D LUTs for the first part of the transformation and a 3x3 matrix for the second. These display profiles, by definition, use the CIE XYZ encoding of the Profile Connection Space. However, using this simpler display profile type does not provide control over the gamut mapping to be performed from PCS colorimetry (simple 1D clipping in XYZ gets applied) and neither does it allow the implementation of perceptual or saturation rendering intents. *N-component LUT-based* display profiles can be used though to provide such control and also to differentiate between the ICC rendering intents.

Transforming scanned data to data for display

Given colorimetric data of the scanned image that is available either using a scanner characterisation model or scanner ICC profile and a characterisation model of the display on which

SCAN TO DISPLAY

it is to be rendered, it is necessary to perform the following transformations, as has been introduced in general previously: First, the colour appearance of the original needs to be computed by taking original viewing conditions (e.g., the viewing conditions in the gallery where the original watercolour is displayed) into account using a colour appearance model. Then, colour and image enhancement may be applied, though in the case of a poster for the exhibition of artwork this is less desirable as the original artwork is of interest as it is. Next, it needs to be ensured that the original's colours are all inside the display's gamut and if they are not then gamut-mapping needs to be applied. Finally the viewing conditions of the display need to be taken into account to determine what stimuli to render on it and the display characterisation model is used to compute digital inputs.

While these steps can be performed individually as described above, in practice they are more likely to be encoded in the scanner and display profiles, whereby the first provides PCS colorimetry of the scan and the second takes that colorimetry and computes digital inputs for the display from it.

The end result of the scenario stages described so far is that if scanner and display characterisation models used take actual viewing conditions into account and if, for example, a simple gamut clipping algorithm is used, then the displayed version of the watercolour will look like the original watercolour (especially as it is unlikely in the case of this media combination that there would be parts of the original that are significantly outside the display gamut). When actual viewing conditions are not taken into account, the relationship between what the profiles represent and what an observer would see can be very weak. To illustrate this point, Figure 7 - 8 shows the measured gamuts of a display and of a reflective original, such as a print or a watercolour, under a range of ambient illumination conditions. Here the profiles may implicitly assume one of the possible states, whereas the actual viewing conditions can be quite different.




Figure 7 - 8: Display (solid line) and reflection print (dashed line) gamuts when viewed simultaneously under varying levels of illumination (CIECAM97s, D50)³⁰.

Editing and page layout

Let's assume that we have successfully scanned our watercolour original, and displayed it on a monitor with an acceptable level of appearance matching between display and original. Our next task is to use the scanned image to create electronic versions of leaflets to be printed on a digital press, and posters to be printed on a large format inkjet printer (Figure 7 - 6). Naturally we want both to display the best possible appearance match to the original, and hence also to each other.

Photo editing programs can be used to adjust scanned images for colour or other properties (scratch & dirt removal, etc.), but we will assume that no such edits are needed for our scanned image. Page layout programs are used to assemble complete page descriptions from their constituting elements, such as raster images (photos), text, lines, and other vector content (typically defined with solid area fill colour). One of the challenges of page composition with respect to colour is the different origin and intent of some of those elements, and the combination of them in a single coherent and good looking whole. In fact this is one of the reasons for the concept of rendering intents to exist in the ICC context. Rendering intents should be associated with *source objects*, not with output devices, and can hence change repeatedly within a single page description. Briefly, the relationship between types of source objects and the "typical" rendering intents they are associated with could be described as follows:

- *Raster images* typically (but not always) originate as photographic captures of a 3D scene or a 2D object, e.g. our watercolour scanned image. Typically the relationships among the pixel colours within the image are more important for the preservation of total image appearance than the individual colour values are, which in general indicates the use of the ICC perceptual rendering intent. In addition, photographic images may represent scenes or objects with colours that are outside of the gamuts of typical hardcopy devices, which would again indicate the ICC perceptual rendering intent. Even though the latter may not be the case for our watercolour image, the former would still indicate the use of the perceptual intent. Some experimentation may be necessary.
- *Renderings* are raster images of a special kind, not originating from the capture of a scene or original, but rather having been generated (rendered) by computer models of objects or scenes. They are increasingly popular for architectural work, for instance, to give potential clients an idea of what buildings would look like in a rather realistic environment. Depending on the sophistication of the computer models used, the results can be almost indistinguishable from captured images of real scenes or objects, or can be quite recognizably "artificial" in shape, texture, lighting, or colour properties. Although much depends on the objects or scenes being modelled and the intended use of the rendered images, in general it is safe to say that renderings typically go with the perceptual rendering intent also.
- *Text* is very different from images in that it is typically defined using primary or secondary (non-half-toned) colours, both for readability and for effect. Especially because of readability, and because of the limitations of many printing technolo-

EDITING AND PAGE LAYOUT

gies, that definition typically goes beyond colour in the strict sense. The well known “black on white” way of rendering text does not only refer to a certain colour of low lightness and chroma, but to the actual inks used to print it: 100% black ink, 0% of all other inks. Although one could typically achieve the same CIE XYZ values with different combinations of printing inks, the result is likely not the same in terms of the sharpness of the character outlines (due to issues such as colour to colour registration). Neither is it likely to be the same in terms of cost, since black ink tends to be cheaper than colour ink. So what rendering intent does this correspond to? Actually *none*, at least in terms of ICC rendering intents. A rendering intent serves to modify a colour conversion down the line, but if an object is being defined using direct device colour definitions, there is no colour conversion or rendering intent involved.

- *Line art*, so called because it is typically made up of straight or curved line segments with the enclosed areas filled with “flat” or “solid” colour, is also known as “vector content” in the technical jargon. In general it behaves more like text than like raster images in terms of colour. For (output) device independent colour specifications the colorimetric or saturation intent is in principle more appropriate than the perceptual intent, but as is the case with text, direct device colour definitions are quite frequent. Named colour (e.g. Pantone[®]) specifications are also frequently used, and represent a challenge of their own. In principle named colours represent certain fixed combinations of colour coordinates in *any* colour space, identified by name. In practice, they often refer to physical samples, often produced with specific ink sets on specific substrates. The translation of named colours into device colour coordinates for display or printing is a complex topic that we will leave aside for now.
- *Logos* could be considered a special case of line art or vector content, often combined with text, sometimes even with images. Since they tend to represent corporate or other identities, requirements for accurate colour reproduction are typically stringent. The use of named colours is common. Most often they would be associated with the colorimetric rendering intent.

A successful page layout application must be able to distinguish between the various object types, and let the user associate different rendering intents with them either on an individual object basis, or on an object type basis, or both. These rendering intents must be preserved in page descriptions generated by the layout application, together with the objects and their corresponding colour definitions. Specific *Page Description Languages* (PDLs) have been developed over the years for this purpose, e.g. PCL, PostScript, or PDF.

We have seen that some objects may be defined using source colour spaces (e.g. scanner RGB) and rendering intents, while others may be defined directly in device colour space coordinates (e.g. text as 100% black ink, 0% all other inks). Both types of definitions must be recorded appropriately in the resulting page descriptions, and must be honoured by the proofing systems, printing systems, or *Raster Image Processors* (RIPs) that interpret them down the line.

Two broadly different approaches can be followed with respect to colour management as controlled by layout applications. The first one, which we might term *greedy colour evaluation*,^{q48} attempts to convert all object colour specifications to output device colour coordinates right away. The result is a PDL file which uses only a single colour space, viz. the output device's, which makes it easy to interpret but by definition device dependent. The gamut mapping and possibly re-rendering transforms that take place during the conversion to device colour coordinates is in general not invertible, so in a sense the document is now *committed* to a specific output device, and that commitment cannot easily be undone. The second approach, which we could term *lazy colour evaluation*, attempts to maintain source object colour specifications as long as possible, without converting to output device colour coordinates. The result is a PDL file, which may use many different colour spaces, which makes it harder to interpret but (output) device independent. Such a file can be *re-purposed* more easily, if necessary.

Let's assume the second strategy for our watercolour reproduction example: both the flyer and the poster compositions will maintain the original object colour specifications to the extent possible.

Proofing

Digital proofing is generally understood as “preparing a sample of printed output on a computer printer before the job is printed on a commercial press” (*Computer Desktop Encyclopedia*⁴⁹). We might extend this definition to include the production of a sample of output of any kind to simulate a sample of output of any other kind. Some commonly distinguished subtypes are:

- *Soft proofing*: a computer display is used to simulate or preview a piece of printed output. Only colour appearance can typically be soft proofed, not other properties such as substrate texture or spatial halftone attributes such as moiré.
- *Hard proofing*: a hardcopy of some kind is used to simulate or preview a hardcopy of a different kind. For instance, a digital inkjet printer can be used to simulate the output of an analogue offset press. All hardcopy proofing systems attempt to simulate colour appearance, and some also attempt to simulate other properties such as substrate texture or spatial halftone attributes.
- *Press proofing*: rather than a simulation this is a sample of actual printing on the intended substrate, produced on the actual printing device. In fact this is not a simulation case but rather a (very) short run length sample of the real thing.

We will limit our discussion here to the simulation of colour appearance, mainly in hard proofing contexts. In general the issues involved with proofing will be easier to resolve, and the simulation more effective, to the extent that the proofing system and the target (final) output system resemble each other. Important dimensions of resemblance include substrates

^q By analogy to *greedy algorithms* from computer science: a type of algorithm that makes the locally optimum choice at each stage with the hope of finding the global optimum, which results in making as many decisions as possible as quickly as possible.

PROOFING

(if any), colorants, marking technology, halftoning characteristics (if any), and last but not least viewing conditions. An extreme example of this is a press proof, where all these characteristics are identical between proofing and target system, even including the viewing conditions. In that case there is obviously no need for any kind of explicit colour management (excepting the usual calibration and process control procedures). It is not too much of an exaggeration to say *that the best colour management is no colour management*, if conditions allow. The more dissimilar the proofing system is to the target system, the more explicit colour management will need to be involved. An extreme example of this is a soft proof, where all the characteristics mentioned are different between proofing and target system, including the viewing conditions. We will now consider the case of (partly) dissimilar proofing and target systems in somewhat greater detail.

Proof printer calibration

Device calibration is essential for repeatable results both from the same printer (*intra-device* consistency) and from different printers (*inter-device* consistency), and the production of digital proof prints is no exception (see Introduction). A typical calibration process consists of the following steps:

- A calibration target is printed;
- The printed target is allowed to dry and/or stabilize completely;
- The printed target is measured with a densitometer, colorimeter, or spectrophotometer;
- Printer calibration parameters are calculated from the measurements;
- Calibration parameters are applied to the printer.

Depending on the calibration algorithm used, the calibration target may consist of density ramps of single colorants or combinations of different colorants. The most traditional calibration algorithms attempt to calibrate both the maximum amount and the tone response of each individual colorant, using optical density based measurements. Since this obviously does not take interactions among colorants into account, more sophisticated algorithms have also been developed that do, typically using colorimetric or even spectral measurements⁵⁰. Which type of algorithm is deemed more suitable for a given proof printer depends in part on the characteristics of its marking engine, in part on characteristics of the hardware and software controlling the printer, and in part also on personal preferences and experience.

Depending on the characteristics of colorants and substrates used, it may be necessary to let printed calibration targets dry or stabilize for a certain amount of time, to make sure that no further density and/or colour changes will occur after measurement. The aim is to calibrate the final appearance of proof prints, not their appearance immediately after printing. Naturally the same criteria should be applied to actual proof prints, in addition to calibration targets.

The calculation and application of calibration parameters naturally depends on the calibration algorithms used. For the traditional maximum density and dot gain (tone response) calibration method, it is normally sufficient to calculate one-dimensional *transfer functions* to be applied to continuous tone colour separated versions of input images (one channel or separation per printing colorant). A transfer function simply maps an input value

COLOUR MANAGEMENT

separation per printing colorant). A transfer function simply maps an input value in a given domain (say [0,1] for simplicity) to an output value in a given range (say [0,1] also) in a deterministic fashion, and without taking any other inputs into account. It should be noted that in order to avoid extrapolation beyond measured data, this type of calibration can only be used to modify the tone response of individual colorants and to *reduce* the maximum amount of colorant, but not to *increase* it beyond the range represented by the calibration target. The latter should therefore represent a larger range than that needed for actual proof prints. To calibrate interactions among colorants a more sophisticated mechanism is needed, for instance multi-dimensional lookup tables.

Proof printer characterization

Device characterization or profiling is essential for colorimetrically accurate results by establishing a correspondence between device-dependent and device independent colour representations, and the production of digital proof prints is no exception. A typical characterization process consists of the following steps:

- A characterization target is printed;
- The printed target is allowed to dry and/or stabilize completely;
- The printed target is measured with a colorimeter or spectrophotometer;
- A printer profile is calculated from the measurements;
- The printer profile is applied in the generation of proof prints.

For most if not all proof printers, characterization targets consist of patches combining different amounts of several colorants, possibly in addition to patches containing single colorants only. The amount and exact kind of the patches used depends on the characterization algorithms used, and in particular on their underlying printer models. Some of the most common printer models are:

- The *Neugebauer* model and its many variants and derivatives; see for instance Mahy⁵¹. The essence of the Neugebauer model is the summation of spectral reflectances of unprinted substrate and solid (full area coverage) single and overprinted colorants (also known as *device states*) to predict the spectral reflectance of non-solid combinations of any number of colorants, weighted by their expected area coverage.
- *Masking equations* – see for instance Berns⁵². Masking equations (whether linear or non-linear) attempt to establish a closed form correspondence between device colour values (e.g. normalized area coverage of individual colorants) and resulting colorimetric or spectral characterizations.
- Lookup tables (LUTs) with interpolation; see for instance Kasson⁵³. Lookup tables are collections of pre-computed (or measured) function values (one- or multi-dimensional) which together with interpolation (linear or non-linear) can be used to estimate the function value for any input within a predetermined domain. As used in printer modelling they are typically 3 or 4 dimensional, mapping for instance CMY or CMYK input values to measured CIEXYZ or CIELAB values, often using multi-linear or tetrahedral interpolation⁵⁴.

PROOFING

In all cases, a *forward* printer model is first established, which maps device colour values to colorimetry. The forward model then gets inverted to produce an *inverse* printer model, which relates colorimetry to device colour values. In proofing contexts both types of model are used: a forward model is first applied to the device colour values of the target device (e.g. an offset press) which results in desired colorimetry values, which then get mapped through an inverse model of the proofing device (e.g. an inkjet printer) to produce the device colour values that when printed will result in the desired colorimetry. The colorimetric proofing process can obviously only work as expected if the proofing device gamut completely encompasses the target device gamut, which we will assume for the sake of our discussion. For more detail on printer characterization models see Bala⁹.

Rendering intents for proofing

As we have seen, proofing is an example of strict colorimetric reproduction (and as such, probably one of the few such examples for practical applications). We have also seen that ICC profiles contain both forward and inverse printer models, usually in the form of lookup tables, in different rendering intent versions. Which rendering intent(s) should we use for proofing then? Given the colorimetric reproduction goal we have two obvious choices: (media) relative colorimetric intent, or absolute (diffuse white relative) intent. The same intent should always be used for both forward and inverse tables.

The most accurate colorimetric reproduction can in principle be obtained with *absolute colorimetric* rendering intents, so shouldn't we just use those? In principle yes, but in practice *your mileage may vary*. Absolute colorimetric reproduction implies that even the colour of unprinted media on the target system will be simulated on the proofing system, resulting in target unprinted media areas being reproduced as printed media areas in the proof print. There are a number of potential issues with this approach:

- The proofing media used must have a lighter (higher) white point than the target media, to accommodate the simulation of the latter. This condition is implicit in the general condition that proofing systems should have a gamut completely encompassing the target system gamut, but it is not always easy to achieve. This is especially true if we want the proofing system to use a media type that has a very similar appearance to the target media, which in general is a very good idea.
- Furthermore, the proofing media should preferably have a white point resembling the target media closely in terms of hue and chroma, to avoid unnatural looking and difficult to achieve and maintain target media white point simulations of a slightly different colour cast.
- To allow proper viewer adaptation to the simulated white point, any non-printed margins showing the original proofing media white point should be removed from proof prints (or from view), otherwise mixed adaptation and very likely unnatural looking simulated target media colour will result.

All of these issues can be avoided or resolved by using the *relative colorimetric* rendering intent instead, which renders target media unprinted areas as proofing media unprinted areas. One would be well advised in this case to look for a proofing media with very similar appearance (in terms of brightness, white point, glossiness, etc.) as the target media, but this is good advice for proofing in general. Any remaining small differences between target and

COLOUR MANAGEMENT

proofing media will normally speaking be ‘absorbed’ by viewer white point adaptation, unless original and proofs are being compared side by side. The latter may be habitual in the evaluation of proofing system capabilities, but not in actual proofing system use – for if one had the original available, there would be little point to producing a simulation of it, to begin with.

It is interesting to compare the digital proofing case, which is what we have been discussing, to the analogue proofing systems of yore. One of the best known examples is no doubt the *3M Matchprint* system, where proofs are “made by exposing the CMYK negatives onto four acetate films which are developed and laminated together”⁵⁵, and on top of a sheet of the actual printing stock that will be used on the target system. The resulting appearance (and colorimetry) will of course be influenced by the properties of the printing stock underneath, and unprinted areas on the target system will show as blank (un-imaged) areas on the proof. Both of these properties are comparable to what is obtained on digital proofing systems with the ICC relative colorimetric intent, and more so than with the ICC absolute rendering intent.

Evaluation of proof prints

In general, proof prints should be evaluated under conditions as similar as possible to the ones that will be used to view or evaluate the final target system prints²². It is worthwhile to note that when ICC profiles are being used for proofing, this in principle implies a single choice of illuminant (D50), illumination level (500 lux), measurement geometry (45/0), standard observer functions (CIE 1931), and so on. If there are significant differences between viewing conditions used for evaluation and measurement conditions used for constructing profiles, a colour match between proof and final print cannot be guaranteed.

To express the differences between a proof and the target printing system that it is meant to proof, colorimetric data can be obtained for a sampling of the target’s gamut and its proofing. The CIEDE2000(1:1) colour difference equation⁵⁶ can then be used to express the differences between corresponding colour pairs and the difference distribution’s statistics can be reported. If the distribution is normal (Gaussian) then parametric statistics are appropriate – mean and standard deviation – and if not, then the median and a high percentile (e.g. 95th) can be reported. In all cases the maximum error needs to be given. However, doing this only expressed how well the proof represents the target’s colours and not how different the colours of a target and proofed images’ will be. To approximate that, it is necessary to take into account factors other than just individual colours and it is more appropriate⁵⁷ to use a high percentile of the difference distribution computed using CIEDE2000(1:2) with optional spatial pre-filtering (<http://www.colour.org/tc8-02/>; MR Luo, pers. comm. 10 March 2006).

Of course there are many non-colour related issues involved with the evaluation of proofs, for instance halftoning properties (moiré effects), media surface finish and appearance, the presence of fluorescence (due to the use of ‘optical brightener’ media additives), media weight, and so on. We have limited our discussion here to the colour related ones.

Challenges and opportunities

We have already hinted at a number of challenges and opportunities involving colour proofing. Here we will briefly outline some more, without going into much detail due to space constraints.

PROOFING

- *Print metamerism.* Metamerism (see Chapter 3) is commonly defined as “the quality of some colours that causes them to appear differently under different light sources. For example, two colour samples might appear the same in natural light, but not in artificial light.”⁵⁸. Besides the light source, the colorants (inks) and substrates being used obviously also affect metamerism. Unless the substrates and colorants of the proofing system are exactly the same as the ones of the target system, they are likely to behave differently to changes in the viewing environment away from the standard viewing conditions. This may cause a proof print to match the target print well only under very constrained viewing conditions that can be difficult to maintain in practice. Even though they may produce matching colorimetric measurements, they can still look different.
- Related to metamerism is also the *colour inconstancy* of a print, which refers to the changes of its appearance when viewed under different light sources. This may cause a print to be an acceptable visual proof under one set of viewing conditions, when viewed on its own (e.g. when sent to a client), and for it to be unacceptable under others. The degree of colour inconstancy can be expressed using the Colour Inconstancy Index – CII⁵⁹.
- *Fading and/or colour stability.* Although proof prints are typically intended for fairly short term use, it is important for them to quickly reach a stable state after printing, and to not be too susceptible to light fade or air fade under normal conditions of use.
- *Proofing actual versus standard devices.* In principle one can generate a colour profile of a target printing device in whatever state it happens to be at a given moment, and use that profile to produce a colour proof. But if the target system is not stable over time or it cannot easily be returned to the condition it was in when the profile was made, the proof will serve little purpose. This observation points to the importance of device calibration for both proofing and target systems, as outlined above. In the commercial printing world this has also led to a movement to target not actual physical printing presses, but a small number of idealized devices that each represent a certain category of real devices and/or applications. The proof is relative to such an ideal device, and the task of actual printing press calibration then becomes to make it resemble the chosen ideal device as closely as possible. In a sense it is then the press that is matching the proof, not the other way around. While that may seem like getting things backwards, in practice it has many advantages and is proving to be a workable approach.

The last point can be rephrased succinctly as “*colour management is not a substitute for process control*”. Process control remains essential for printing in general, and includes such essentials as device calibration. It is not because we *can* describe arbitrary states of printing systems in a colour profile that we *should*.

Poster and leaflet production

We have now almost come to the end of our watercolour reproduction scenario. We have captured an original, examined and/or edited it on a computer and electronic display, and have produced colour proofs of leaflets to be printed on a digital press and posters to be printed on a large format inkjet printer. Now all that is left for us to do is print the leaflets and posters, distribute them all over town, and organize the opening reception for our art exhibition.

During our discussion of proofing above we have finessed a subtle but important point. Once the device colour values of a target system are known, we can use them together with a colour profile of that system to derive the corresponding colorimetry, which we can then use to generate a colour proof. But how do we determine the desired device colour values for the different systems involved? Assuming that we have calibrated each device correctly, and that we have accurate colour profiles available for each, we can map the desired colorimetry through each profile to obtain the corresponding device colour values. So what is the desired colorimetry? If we have colorimetrically accurate capture data available, we can simply choose that. This would basically bring us back to the colorimetric proofing scenario, except that now we are proofing original colorimetry (*artist colorimetry*, as it were) rather than some target printing system colorimetry. As we have discussed, proofing implies using a colorimetric rendering intent, whether relative or absolute. But as we also discussed, colorimetric proofing really only makes sense under very constrained circumstances, including viewing environments and substrate types, which we most likely cannot maintain for our leaflet and poster production scenario.

If it is not colorimetric proofing that we are after, then what is it? Since it would lead us too far to discuss the issues involved in detail, suffice it to say that we are striving for reproductions that are adapted to or optimised for each printing system and substrate type involved, while still remaining faithful to the *look and feel* of the original. In terms of ICC colour management this indicates the use of perceptual rendering intents rather than colorimetric ones. In terms of print evaluation this suggests that each print should be taken on its own terms (and in its own viewing environment), while still keeping an eye on the original. A tall order for which unfortunately there are no well-defined recipes that will guarantee us the results we are after. In the next and final section of this chapter we will discuss some future opportunities that may help us achieve our goal.

Future opportunities

No design is perfect, and although the current CIE-based colour management systems can reach excellent levels of performance and sophistication, a number of opportunities for future improvement can be identified. We will just mention a few of these, sketching the expected benefit and indicating how they are related to past or ongoing work within the CIE.

Self-calibrating and self-profiling devices

Device calibration (“to adjust or bring into balance”⁶⁰), is fundamental to reliable colour management, as we have indicated above. Calibration brings a device into a known state, without which device characterization data as represented in a device profile is meaningless. Nearly all colour capture, display, and reproduction devices in use today have considerable analogue subsystems, which tend to require periodic calibration to compensate for drift. To the extent that calibration procedures can be made automatic and unobtrusive, the reliability and quality of colour management systems will increase, as will their user friendliness. Most modern desktop scanners have automatic calibration capabilities, for instance using internal white reference strips. Displays can be calibrated, but typically not automatically and unobtrusively. Some high-end displays come with dedicated hardware (sensors) and software for calibration, but the onus is on the user to periodically take out the sensor, mount it on the display, and trigger the calibration procedure. To make progress, calibration sensors should be incorporated into the display itself (using wave guides or other appropriate means to measure part of the emitted light), and calibration procedures should be triggered and run their course without the need for user intervention (for instance, in the disguise of a “screen saver” that is run during computer idle time). Recently a number of inkjet and other printers have appeared with built-in calibration sensors (typically densitometers or derivatives), with the corresponding automatic calibration routines (as so-called “firmware”). Even though such devices can automatically trigger calibration routines during idle time or when the necessary conditions are met, the calibration procedures themselves can hardly be called unobtrusive. They typically involve the printing of a test pattern on the available media, its measurement using the built-in sensor, and the calculation and application of a set of calibration parameters. Nevertheless, this is major progress compared to using off-line measurement and calibration procedures.

Once a device is calibrated, it is in a known state. Exactly what that state looks like needs to be described by collecting device characterization data and constructing a device model from it, typically in the form of a device profile. For capture devices this involves capturing a profiling target with known device-independent colour specifications (using one of the CIE colour spaces), and relating the latter to the device colour values resulting from the capture. Conversely, for display and reproduction devices this involves (re)producing a profiling target with known device colour values, measuring the corresponding device output in device-independent (CIE based) colour coordinates, and relating the two. Both types of profiling procedures could in principle be made automatic. For scanners one could imagine a colour characterization target being built into the device itself, similar to the white references being used today. This, however, would not take into account the effect of different types of “original” media being scanned, which is necessary for best performance. For displays, the same type of “screen saver” approach could be used as described above for calibration, but the built-in measurement device may have to be upgraded to a colorimeter or spectrophotometer that can be traced to CIE standards (which is not necessarily required for calibration only). Likewise, printers would need a built-in colorimeter or spectrophotometer traceable to CIE standards, but otherwise the process would be comparable to the ones in use today for automatic calibration.

Workflow automation

In technology, *workflow* typically means “the automatic routing of documents to the users responsible for working on them”⁶¹. With respect to colour management, we could perhaps rephrase this as *the routing of colour content to the system components responsible for transforming it*. We can think of devices such as digital cameras, scanners, displays, and printers, and software for capturing, editing, and preparing for (re)production as some of these system components. If such transformations of colour content are to be automatic, something desirable in all but the most high end of application contexts, then it must be represented in an unambiguous way, and certain *division of labour* agreements must be in place among the system players.

To illustrate the first requirement, it is not enough to know that a certain file contains RGB data for it to be displayed correctly. At the very least it must be made clear what kind of RGB is involved, for instance by relating it to a CIE based device characterization profile of the device that produced it. Likewise, it is not sufficient to know that a document contains CMYK data for it to be printed correctly. Reference could be made to another CIE based device characterization profile of the device that the document was *intended* to be reproduced on. Such links between colour encodings and device characterization profiles can be made via the embedding of the latter in digital file formats, or via other mechanisms such as metadata tags referring to a set of standardized profiles. Unless all data in a given system shares the same image state and expected viewing conditions²⁷, such properties must also be made explicit, and system components must be prepared to deal with them.

Unambiguous representation of exchanged colour data is necessary, but not sufficient for everything to work as expected. One of the most common causes of problems in colour managed systems is the proliferation and redundancy of colour management functionality among different devices and software components of heterogeneous systems^f. This often leads to unintended *double colour management* (or worse) and poor results, which in turn needs complex configurations of each component involved to try and prevent it. Figure 7 - 9 provides a schematic illustration of the issue.

To avoid these problems, one needs to know what the user intends to do with a certain document, which types of data or colour encodings each system component is capable of accepting, what types of transformations it can apply to them, exactly which transformations have been applied at each step along the way, and hence which types of transformations remain to be applied before the document is ready for (re)production at its intended destination. These things can be achieved by careful manual configuration of each system component, but in more complex systems that eventually requires a colour scientist to be put in each product box (and they don't come *that* cheap). The alternative is automatic configuration using a kind of *universal plug and play*⁶² or *zeroconf* approach⁶³. As mentioned above, colour needs to be managed, but that shouldn't require an MBA or PhD.

Automatic adaptation to viewing environment

One of the limitations of earlier CIE colour spaces such as CIE XYZ or CIELAB was that they could not deal easily with different viewing environments, implicitly being defined only

^f In the sense of components (hardware or software) from different manufacturers being involved.

FUTURE OPPORTUNITIES

for one particular, and fixed, such environment. The newer generations of CIE colour appearance spaces such as CIECAM97 and CIECAM02 do take viewing conditions into account explicitly and in a parameterized fashion. What this means in practice is that if one knows the intended (or actual) viewing environment for a particular document, one can process its colour data such that the result will appear “correct” in that environment. But where does one get the required information about the viewing environment? In some cases, for instance the sRGB systems discussed above, one can just assume that it is fixed, and hope for the best. In other cases, for instance in the current ICC based systems, the implicit viewing environment parameters are normative rather than informative, and hence the results produced by such systems are only “valid” if actual viewing conditions match the prescribed ones. Unfortunately such constraints cannot be successfully imposed in many practical scenarios, and usually it is not much of a consolation to tell an unhappy user that “your viewing environment is wrong, hence your complaint is not admissible”.




Figure 7 - 9: Proliferation and redundancy of colour management functionality in heterogeneous colour managed systems. The large shaded boxes represent RGB based applications, operating system functionality, CMYK based applications, raster printer drivers, and printer embedded PS/PDF RIPs respectively, in English reading order. Smaller boxes inside the former represent colour profile and colour data types, and lines among them the different ways in which they can in principle be connected. Although in principle there is only one colour transformation from input to output colour space involved, a system as depicted allows 27 different paths to be constructed. Almost half of those result in wrong output, and about 75% of them are redundant, resulting in the same output as some other one.

COLOUR MANAGEMENT

If viewing condition parameters could be established (measured) automatically, perhaps the necessary adaptations could also happen automatically and transparently to the user? In the case of digital cameras, the viewing environment is not independent of the scene being captured, in fact in a sense it *is* the scene being captured. Most modern digital cameras do record a number of types of metadata together with the image, such as time (and sometimes place) of capture, aperture and shutter speeds, etc. Perhaps this would need to be extended to more scene-related parameters such as brightness level, estimated illuminant colour temperature or spectral characteristics, etc. For scanners the situation is rather different, since they self-illuminate an original (whether reflective or transmissive) with a fixed light source using a fixed illumination geometry. A description (perhaps added to a scanner profile) of the characteristics of the light source and illumination geometry used might help in some cases, but the utility of this seems limited. Perhaps an “active scanner” concept would be more to the point: rather than using a single fixed “viewing environment”, such a device would accept instructions to “view” an original being scanned under the conditions that the image requester would consider appropriate. One could think of brightness level, illumination geometry, and perhaps even colour temperature or spectral composition of the light source used for scanning.⁵

Electronic display or “softcopy” devices are unique when compared to input and hardcopy devices, in the sense that the images produced by them cannot be “transported” to different viewing environments than the one they were produced in.¹ The images produced are ephemeral and inextricably linked to the device itself (*the medium is the message*, quite literally perhaps). As such there is no need to worry about recording viewing environment parameters for later use, only for instantaneous use. As an example of this approach, a well-known colour management hardware and software manufacturer has recently introduced a consumer level monitor calibration system that has an interesting and unusual feature. The sensor can be put into a small stand and left on the desk next to the monitor (i.e. in the same viewing environment as the monitor), and under software control will periodically measure the ambient light level and adjust the monitor’s brightness to compensate for any changes. Although the idea is not new (some TV sets already used a “magic eye” for much the same purpose in the 60’s, and some clock radios include a similar circuit for dimming the display brightness at night), it is a nice example of automatic measurement of and adaptation to viewing environment parameters in the digital colour management domain. Brightness control is probably the easiest thing to do, but perhaps more sophisticated types of measurement and adaptation will follow.

Hardcopy devices (printers) share with scanners and digital cameras that the images they produce can and typically are “transported” to other viewing environments than the one they have been produced in. Hence the measurement of and adaptation to production viewing environments would only serve a temporary purpose, and might actually make things worse for the “end user” of the images produced. Mirroring the suggested capability of scanners to “view things a certain way”, one could imagine printers being capable of “producing images

⁵ Many spectrophotometers use a primitive version of this already, in the form of a UV cut filter that can be placed into or removed from the illumination path.

¹ Ignoring corner cases like electronic displays being photographed and printed on paper, for the sake of the argument.

FUTURE OPPORTUNITIES

a certain way”, to make them suitable for a particular viewing environment. Some of this is possible today, for instance via the use of different printer profiles intended for different viewing environments, but this type of control is rather coarse, discrete, and clumsy in use. Other options might exist, for instance changing the way colour separations are being calculated inside (or for) the printer to optimise them for certain intended viewing environments. An example might be the reduction of illuminant metamerism, or the increase of colour constancy, for a certain set of illuminant types. The most principled solution for the latter is more likely to be truly spectral image reproduction (see below), but lacking that, there may be other possibilities.

Spatial processing

At present the operations needed for managing colour are typically performed on a colour-by-colour basis, which means that all instances of an original colour are reproduced as the same reproduction colour. While this is certainly a reasonable starting point, benefits can be had from taking a colour’s spatial neighbourhood in the original into account when determining its reproduction. Examples of such spatial approaches are spatial gamut mapping algorithms^{64,65} as well as efforts to model colour appearance of the parts of spatial complex stimuli, such as photographic images (see the iCAM model⁶⁶). The CIE’s TC8–08 on *Spatial Appearance Models* is also active in this area.⁶⁷

Smart CMMs

CMMs are responsible for transforming colour data from an input colour space to an output colour space, using input and output colour profiles and additional parameters such as rendering intent. In the current ICC architecture, the role of the CMM is essentially limited to that of an interpolation engine, whereas the colour profiles contain all the “value add” such as gamut mapping, re-rendering transforms, and perhaps even viewing condition compensation. All of these things are “built in” at profile creation time, and cannot be changed afterwards. This type of arrangement has been described as a “smart profile, dumb engine” architecture. The opposite of this would then be a “dumb profile, smart engine” architecture, which might have certain advantages over the former:

- Device profiles would contain little more than colorimetric (or spectral) measurements, which would make them very easy to produce and liberate them of undocumented and vendor specific (incompatible and non-portable) “secret sauce”.
- Viewing conditions could be specified explicitly and independently of device profiles, and taken into account by the smart engine when calculating colour transforms.
- Gamut mapping and/or re-rendering transforms could be specified explicitly and independently of device profiles, and applied with knowledge of the specific pair of devices or even images (and their respective gamuts) for which a colour transform is being calculated.

Such a system would seem more flexible than the current one, with the potential for improved results especially considering heterogeneous systems and/or often or quickly changing device pairings. Also things like automatic adaptation to viewing environments, as discussed above, could in principle be accommodated more easily. So what are we waiting

COLOUR MANAGEMENT

for? As always, there are also a number of potential drawbacks or difficulties with such a radically different architecture:

- The “secret sauce” that would be eliminated from device profiles is actually there for a good reason. Calculating gamut mappings and re-rendering transforms is usually not an easy task, and perhaps one that cannot be accomplished with sufficient quality in real time (or at least, an acceptably short time for a smart CMM scenario).
- It is not immediately evident how to specify viewing conditions in sufficient detail, or how to use such specifications in on–the–fly construction of colour transformations between arbitrary pairs of devices, or colour spaces.
- Gamut mapping is in fact a big part of the aforementioned “secret sauce”, hence the same considerations apply to it. While it is certainly conceptually interesting to be able to separate gamut mapping algorithms from the data (and colour spaces) they operate on, it is not evident how such a separation could be implemented in practice, and with sufficient quality and speed.

Although the concept of smart CMMs has been around for quite a while, until recently little progress had been made towards its realization in practical architectures and implementations. The recent announcement of Microsoft’s *Windows Color System* (WCS)⁶⁸ has changed that. Scheduled for introduction in 2006, together with Microsoft’s next generation operating system dubbed *Vista*, WCS effectively aims to implement a smart CMM colour management system. The issues mentioned above are addressed in the following way:

- *Device profiles* are indeed simple XML–formatted “containers” of colorimetric device measurements, specified in CIE XYZ and related colour spaces. *Device models* are responsible for turning those measurements into complete forward (and inverse) models of the devices at hand. Some types of device models are provided as a standard part of WCS, others can be provided as *plug–ins* by device manufacturers or other 3rd parties.
- Viewing conditions are separate from device models, and are specified using CIECAM02.
- Gamut mapping is separate from device models and viewing conditions. Some gamut mapping algorithms are provided as a standard part of WCS, and these are derived from the work of CIE TC 8-08 on gamut mapping. Other algorithms can be provided as plug–ins by device manufacturers or other 3rd parties, much as is the case for device profiles.

It is not our intention here to evaluate the quality or potential of the WCS system, but merely to flag its existence as a potentially important step in the evolution of practical colour management systems. Another worthwhile observation is that although WCS represents a fairly radical departure from existing systems, it is nevertheless clearly and explicitly based on the work of the CIE in its various shapes and forms.

It could be expected that the announcement of WCS might rekindle interest in smart CMM architectures within the ICC, where much discussion and even some prototyping has taken place on this subject in the past.

Multi-spectral imaging (CIE TC8-07)

As more and more imaging devices are capable of more than trichromatic colour reproduction, the possibility arises for reproducing an original's spectral properties (as opposed to only the tristimulus) (see Chapter 1 of Hunt¹). For self-luminous originals this would give a match in the spectral power distribution of original and reproduction and as a consequence the two would look identical to observers with normal colour vision, to those who have deficient colour vision and also to those animals whose visual systems are sensitive to electromagnetic radiation in the range where the spectral match holds. For reflective and transmissive originals a spectral match would be in terms of spectral reflectance or transmittance and, in addition to the properties of spectral matches of self-luminous originals, it would result in the original and reproduction looking the same under any illumination (i.e. their appearance would change in the same way for each change of light source).

The key challenges in digital multi-spectral colour reproduction include questions about the encoding of multi-spectral data (e.g. Uchiyama *et al.*⁶⁹), the spectral characterization of imaging devices (e.g., Chen *et al.*⁷⁰) and the potential benefits to gamut mapping from working in a spectral domain (e.g. Derhak and Rosen⁷¹). The CIE's TC 8-07 on *Multispectral Imaging* (<http://www.colour.org/tc8-07/>) too is active in this field.

Conclusion

We hope to have provided the reader with a reasonable overview of colour management systems past, present, and future, and their importance to an ever increasing number of applications. If anything is clear from our discussion, it is that the work of the CIE has been an essential part of most if not all of these developments. We sincerely hope that this will continue to be the case for the foreseeable future.

Acknowledgements

Eugenio Martinez-Uriegas (Hewlett-Packard Company) for the idea of representing colour comparisons as comparisons between observer-internal entities as used in Figure 7 - 1; Craig Revie (Fuji Film Electronic Imaging Ltd.), Eric Walowit (Color Savvy Systems Inc.) and Jack Holm (Hewlett-Packard Company) for advice on future developments within the ICC.

References

- 1 Hunt RWG (1995) *The Reproduction of Colour*, 5th Edition, Fountain Press, England.
- 2 CIE Technical Report (2004) Colorimetry, 3rd ed. Publication 15:2004, CIE Central Bureau, Vienna.
- 3 CIE Technical Report (2004): A colour appearance model for colour management systems: CIECAM02 Publication CIE 159:2004 CIE Central Bureau, Vienna.
- 4 Yendrikhovskij SN (1998) *Color Reproduction and the Naturalness Constraint*, Ph.D. thesis, Technische Universiteit Eindhoven, The Netherlands, ISBN 90-386-0719-9.
- 5 de Ridder H (1995) Naturalness and Image Quality: Influence of Chroma Variation at Various Lightness Levels, *IS&T/SID 3rd Color Imaging Conference*, Scottsdale, AZ, 140-143.

- 6 Xiao K, Li C, Luo MR, Taylor C (2004) Colour Appearance for Dissimilar Size, *IS&T 2nd European Conference on Color in Graphics, Imaging and Vision (CGIV)*, Aachen, Germany, 65–69.
- 7 de Fez MD, Luque JM, Capilla P, Pérez–Carpinell J, Díez MA (1998) Colour memory matching analysed using different representation spaces, *J. Optics*, IOP Publishing Ltd., England, **29**:287–297.
- 8 IEC (1999) IEC 61966–2–1: Multimedia systems and equipment – Colour measurement and management – Part 2–1: Colour management – Default RGB colour space – sRGB, IEC, Switzerland.
- 9 Bala R (2003) Device characterisation, *Digital Color Imaging Handbook*, Sharma G. (ed.), CRC Press, 269–384.
- 10 Fairchild MD, Johnson GM (2002) Meet iCAM: A Next–Generation Color Appearance Model, *IS&T/SID 10th Color Imaging Conference*, Scottsdale, AZ, 33–38.
- 11 CIE Technical Report (2004) Chromatic Adaptation Under Mixed Illumination Condition When Comparing Softcopy and Hardcopy Images. Publication CIE 162:2004 CIE Central Bureau, Vienna.
- 12 Fairchild MD (2004) *Color Appearance Models*, 2nd Edition, John Wiley and Sons.
- 13 Johnson A (1992) Colour Appearance Research for Interactive System Management and Application – CARISMA, Work Package 2 – Device Characterisation, Report WP2–19 Colour Gamut Compression (For a description of the CARISMA algorithm see also (Morovič, 1999, pp. 271–274)).
- 14 Cholewo TJ, Love S (1999) Gamut Boundary Determination Using Alpha–Shapes, *IS&T/SID 7th Color Imaging Conference*, Scottsdale, AZ, 200–204.
- 15 Braun GJ, Fairchild MD (1997) Techniques for Gamut Surface Definition and Visualization, *IS&T/SID 5th Color Imaging Conference*, Scottsdale, AZ, 147–152.
- 16 Morovič J, Luo MR (1999) Developing Algorithms for Universal Colour Gamut Mapping. In *Colour Engineering: Vision and Technology*, MacDonald LW, Luo MR (eds.), John Wiley & Sons, Chichester, England, 253–282.
- 17 Morovič J, Luo MR (2000) Calculating Medium and Image Gamut Boundaries for Gamut Mapping, *Color Research and Application*, **25**:394–401.
- 18 Morovič J (2003) Gamut Mapping, *Digital Color Imaging Handbook*, Sharma G. (ed.), CRC Press, 639–685.
- 19 CIE (2004) Guidelines for the Evaluation of Gamut Mapping Algorithms, Publication 156:2004, CIE Central Bureau, Vienna.
- 20 Stokes M (1997) Industry Adoption of Color Management Systems, *Proceedings of AIC Color 97 Kyoto*, 126–131.
- 21 ICC (2004) Specification ICC.1:2004–10 (Profile version 4.2.0.0) Image technology colour management – Architecture, profile format, and data structure, <http://www.color.org/ICC>.
- 22 ISO (2000) ISO 3664:2000. Viewing conditions – Prints, transparencies and substrates for graphic arts technology and photography, ISO.
- 23 Borg L (2004) *Black Point Compensation From Adobe Systems*, Adobe Systems, <http://www.color.org/Adobe1bpc.pdf> [accessed 15 November 2005].

REFERENCES

- 24 Koh KW, Tastl I, Nielsen M, Berfanger DM, Zeng H, Holm J (2003) Issues Encountered in Creating a Version 4 ICC sRGB Profile, *IS&T/SID 11th Color Imaging Conference*, Scottsdale, AZ, 232–237.
- 25 Adobe Systems Inc. (2005) *Adobe® RGB (1998) Color Image Encoding*, Version 2005–05, Adobe Systems Inc.
- 26 Merriam–Webster Inc. (2003) *Merriam–Webster’s 11th Collegiate Dictionary*, Merriam–Webster Inc.
- 27 ISO (2004) ISO 22028–1:2004. Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange — Part 1: Architecture and requirements, ISO.
- 28 ISO (1997) ISO 12641:1997 Graphic technology – Prepress digital data exchange – Colour targets for input scanner calibration, ISO.
- 29 Morovič PM (2002) *Metamer Sets*, PhD Thesis, School of Information Systems, University of East Anglia.
- 30 Morovič J, Sun PL, Morovič PM (2001) The gamuts of input and output colour reproduction media, *Proc. SPIE*, **4300**:114–125.
- 31 Schroedinger E (1920) Theorie der Pigmente von groesster Leuchtkraft, *Ann. Phys. (Paris)*, **62**:603–622.
- 32 Morovič J, Morovič PM (2003) Determining Colour Gamuts of Digital Cameras and Scanners, *Col. Res. & Appl.*, **28**/1 59–68.
- 33 Finlayson GD, Morovič PM (1999) Metamer Constrained Color Correction, *IS&T/SID 7th Color Imaging Conference*, Scottsdale, AZ, 26–31.
- 34 Yule JAC (1938) The Theory of Subtractive Colour Photography. I. The Conditions for Perfect Colour Rendering, *J. Optical Society of America*, **28**:419–430.
- 35 Kang HR (1992) Color Scanner Calibration, *J. Imaging Sci. Tech.*, **36**/2:162–170.
- 36 Kang HR, Anderson PG (1992) Neural Network Applications to the Color Scanner and Printer Calibrations, *J. Electronic Imaging*, **1**/2:125–135.
- 37 Gatt A, Morovic J, Noriega L (2003) Colorimetric Characterization of Negative Film for Digital Cinema Post–Production, *IS&T/SID 11th Color Imaging Conference*, Scottsdale, AZ, 341–345.
- 38 Kasson JM, Nin SI, Plouffe W, Hafner JL (1995) Performing Color Space Conversions with Three–Dimensional Linear Interpolation, *J. Electronic Imaging*, **4**/3:226–250
- 39 Shepard D (1968) A two–dimensional interpolation function for irregularly spaced data, *Proc. ACM*, ACM Press, 517–524.
- 40 Spaulding, KE, Woolfe GJ, Giorgianni EJ (2000) Reference Input/Output Medium Metric RGB Color Encodings (RIMM/ROMM RGB). *IS&T PICS 2000 Conference Proceedings*, Portland, OR, 155–163.
- 41 IEC (2003) IEC 61966–2–2: Multimedia systems and equipment – Colour measurement and management – Part 2–2: Colour management – Extended RGB colour space – scRGB, IEC, Switzerland.
- 42 Berns RS, Motta RJ, Gorzynski ME (1993a) CRT Colorimetry. Part I: Theory and Practice, *Color Research and Application*, **18**/5:299–314.

- 43 Berns RS, Gorzynski ME, Motta RJ (1993b) CRT Colorimetry. Part II: Metrology, *Color Research and Application*, **18**/5:315–325.
- 44 Kwak Y, MacDonald LW (2000) Characterisation of a desktop LCD projector, *Displays*, **21**/5:179–194.
- 45 Post DL, Calhoun CS (1989) An Evaluation of Methods for Producing Desired Colors on CRT Monitors, *Color Research and Application*, **14**:172–186.
- 46 Murakami Y, Hatano N, Takiue J, Yamaguchi M, Ohyama N (2004) Evaluation of smooth tonal change reproduction on multiprimary display: comparison of color conversion algorithms, *Proc. SPIE*, **5289**:275–283.
- 47 CIE Technical Report (1996) The Relationship between Digital and Colorimetric Data for Computer–Controlled CRT Displays, Publication 122:1996. CIE Central Bureau, Vienna.
- 48 Greedy algorithm." Wikipedia. Wikipedia, 2005. *Answers.com* 21 Dec. 2005. <http://today.answers.com/topic/greedy-algorithm>.
- 49 "Digital proofing." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 06 Mar. 2006. <http://today.answers.com/topic/digital-proofing>.
- 50 Bala R, Sharma G, Monga V, Van-de-Capelle JP (2005) Two–dimensional transforms for device color correction and calibration, *IEEE Trans. Image Proc.*, **14**/8:1172–1186.
- 51 Mahy M (1997) Calculation of color gamuts based on the Neugebauer model, *Color Research and Application*, **22**/6:365–374.
- 52 Berns RS (1997) A generic approach to color modeling, *Color Research and Application*, **22**/5:318–325.
- 53 Kasson JM, Plouffe W (1992) An Analysis of Selected Computer Interchange Color Spaces, *ACM Transactions on Graphics*, **11**/4:373–405.
- 54 Hardeberg JY, Schmitt F (1997), Color Printer Characterization Using a Computational Geometry Approach, *IS&T/SID 5th Color Imaging Conference*, Scottsdale, AZ.
- 55 "Matchprint." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 09 Mar. 2006. <http://today.answers.com/topic/matchprint>.
- 56 CIE Technical Report (2001) Improvement to industrial colour difference evaluation. Publication 142–2001 CIE Central Bureau, Vienna.
- 57 Uroz J, Morovic J, Luo M R (2002) Perceptibility Thresholds of Colour Differences in Large Printed Images, in *Colour Image Science: Exploiting Digital Media*, MacDonald L. W. and Luo M. R. (eds.), John Wiley & Sons, 49–73.
- 58 "Metamerism." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 09 Mar. 2006. <http://today.answers.com/topic/metamerism>
- 59 Luo MR, Li CJ, Hunt RWG, Rigg B and Smith KJ (2003) The CMC 2002 Colour Inconstancy Index: CMCCON02, *Coloration Technology*, **119**:280–285.
- 60 "Calibrate." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 17 Dec. 2005. <http://today.answers.com/topic/calibrate>
- 61 "Workflow." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 17 Dec. 2005. <http://today.answers.com/topic/workflow>

REFERENCES

- 62 "Universal plug-and-play." Computer Desktop Encyclopedia. Computer Language Company Inc., 2005. *Answers.com* 17 Dec. 2005. <http://today.answers.com/topic/universal-plug-and-play>
- 63 "Zeroconf." Wikipedia. Wikipedia, 2005. *Answers.com* 17 Dec. 2005. <http://today.answers.com/topic/zeroconf>
- 64 Bala R, DeQueiroz R, Eschbach R, Wu W (2001) Gamut mapping to preserve spatial luminance variations, *Journal of Imaging Science & Technology*, **45/5**:436–443.
- 65 Morovič J, Wang Y (2003) A Multi-Resolution, Full-Colour Spatial Gamut Mapping Algorithm, *IS&T/SID 11th Color Imaging Conference*, 282–287.
- 66 Fairchild MD, Johnson GM (2004) The iCAM framework for Image appearance, image difference, and image quality, *Journal of Electronic Imaging*, **13**:126–138.
- 67 Johnson GM (2005) Cares and Concerns of CIE TC8–08: Spatial Appearance Modeling & HDR imaging, *SPIE/IS&T Electronic Imaging Conference*, San Jose, 148–156.
- 68 Microsoft (2005), Windows Color System: The Next Generation Color Management System, Microsoft Corporation.
- 69 Uchiyama T, Yamaguchi M, Haneishi H, Ohyama N, Nambu S (2004) A Method for the Unified Representation of Multispectral Images with Different Number of Bands, *Journal of Imaging Science and Technology*, **48/2**:120–124.
- 70 Chen Y, Berns RS, Taplin LA, Imai FH (2003) A Multi-Ink Color-Separation Algorithm Maximizing Color Constancy, *IS&T/SID 11th Color Imaging Conference*, Scottsdale, AZ, 277–281.
- 71 Derhak MW, Rosen MR (2004) Spectral Colorimetry Using LabPQR – An Interim Connection Space, *IS&T/SID 12th Color Imaging Conference*, Scottsdale, AZ, 246–250.