

Standard RGB Color Spaces

Sabine Süsstrunk

Laboratory of audio-visual Communication (EPFL), Lausanne, Switzerland

Robert Buckley

Xerox Architecture Center, Webster, New York, USA

Steve Swen

Apple Computer, Cupertino, California, USA

Abstract

This paper describes the specifications and usage of standard RGB color spaces promoted today by standard bodies and/or the imaging industry. As in the past, most of the new standard RGB color spaces were developed for specific imaging workflow and applications. They are used as interchange spaces to communicate color and/or as working spaces in imaging applications. Standard color spaces can facilitate color communication: if an image is in 'knownRGB,' the user, application, and/or device can unambiguously understand the color of the image, and further color manage from there if necessary. When applied correctly, a standard RGB space can minimize color space conversions in an imaging workflow, improve image reproducibility, and facilitate accountability.

The digital image color workflow is examined with emphasis on when an RGB color space is appropriate, and when to apply color management by profile. An RGB space is "standard" because either it is defined in an official standards document (a *de jure* standard) or it is supported by commonly used tools (a *de facto* standard). Examples of standard RGB color spaces are ISO RGB, sRGB, ROMM RGB, Adobe RGB 98, Apple RGB, and video RGB spaces (NTSC, EBU, ITU-R BT.709). As there is no one RGB color space that is suitable for all imaging needs, factors to consider when choosing an RGB color space are discussed.

Introduction

With the democratization of digital imaging technology, managing color in an unambiguous way has become one of the major concerns of the imaging industry. If the image doesn't "look good," consumers will not invest in digital imaging. How to explain to consumers why their images look different on their monitors and their relatives' monitors, not to mention their 'photographic' printers? "Good color" is not just the concern of color scientists and engineers anymore, it has become a marketing goal.

The professional imaging market is concerned with communicating color in unambiguous ways to an even greater extent. Closed production environments where skilled operators manage color are becoming a thing of the past. It is quite common today for an image to be scanned

by a digital photo agency in Europe, integrated into a publication layout in the United States, and printed in Asia.

The ICC architecture, managing color by tagging an image with a profile that contains information about its color, has been successfully implemented in some but not all aspects of the imaging workflow. The current ICC architecture is not unambiguous enough for many professional applications, and not transparent enough for many consumer applications.¹ There are still applications and devices that are not ICC compatible. Managing color with ICC profiles does not always result in predictable reproductions. If an image looks bad, who is at fault? The 'bad' scan, the 'bad' profile, the 'bad' CMM, or the 'bad' workflow?

The need for good color by professionals and consumers alike has resulted in the recent development of new, "standard" RGB color spaces, promoted by the imaging industry and sometimes adopted by standard bodies. People are familiar with RGB: their scanners capture RGB, their applications work in RGB channels, and their monitors display with RGB primaries. However, these RGB spaces are based on different spectral attributes and conversions to and from these spaces are still necessary. Before examining the characteristics of RGB color spaces and discussing their usage, it is therefore necessary to look at the digital image color workflow as a whole.

Digital Image Color Workflow

According to the CIE, a color space is a geometric representation of colors in space, usually of three dimensions.² The basis functions are color matching functions, usually CIE color matching functions. Spectral spaces are spaces spanned by a set of spectral basis functions. The set of color spaces is therefore a subset of the set of spectral spaces. However, in practice, the difference is often neglected, and all representations of color in space are called a "color space."

The color flow of a digital image can be generalized as follows.³ An image is captured into a sensor or source device space, which is device and image specific. It may then be transformed into an unrendered image space, i.e. a standard color space describing the original's colorimetry. In most workflows, however, the image is directly

transformed from the source device space into a rendered image space, which describes the color space of some real or virtual output. If the rendered image space describes a virtual output, then additional transforms are necessary to convert the image into an output space, which is an output device specific color space.

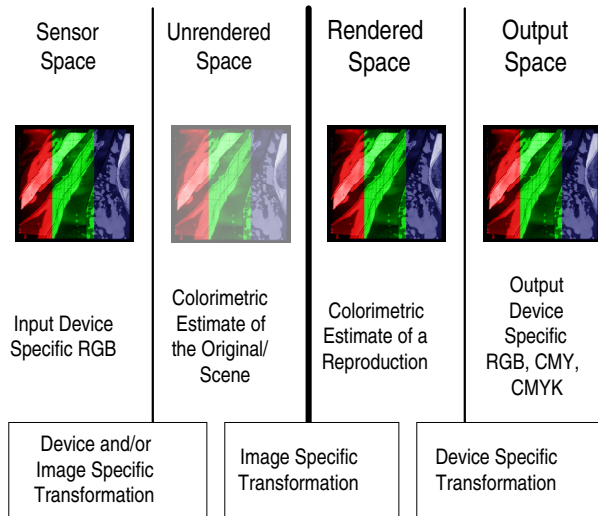


Figure 1. Schematic representation of a digital image color flow

Standard RGB color spaces will always describe either unrendered or rendered image spaces; most existing standard RGB color spaces fall into the category of rendered image spaces. Source and output spaces are always device specific.

Sensor Space

When a scene or original is captured, either by a scanner or by a digital camera, its first color space representation is device and scene specific, defined by illumination, sensor, and filters (Figure 2). In the case of scanners, the illumination should be constant for each image. With digital cameras, the illumination can vary from scene to scene, and even within a scene. A source specific RGB is not a CIE-based color space, but a spectral space defined by the spectral sensitivities of the camera or scanner.

When images are archived or communicated in sensor space, camera or scanner characterization data, such as device spectral sensitivities, illumination, and linearization data have to be maintained so that further color and image processing is possible.^{4,5} Ideally, the image should be saved in a standard file format, such as TIFF/EP, which has defined tags for the necessary information.⁶

It is highly unlikely that there will ever be a “standard” source RGB space. With digital cameras, the illumination is scene dependent. With scanners, manufacturers would have to agree on using the same light source, sensors, and filters—components that are typically selected on the basis of engineering considerations.

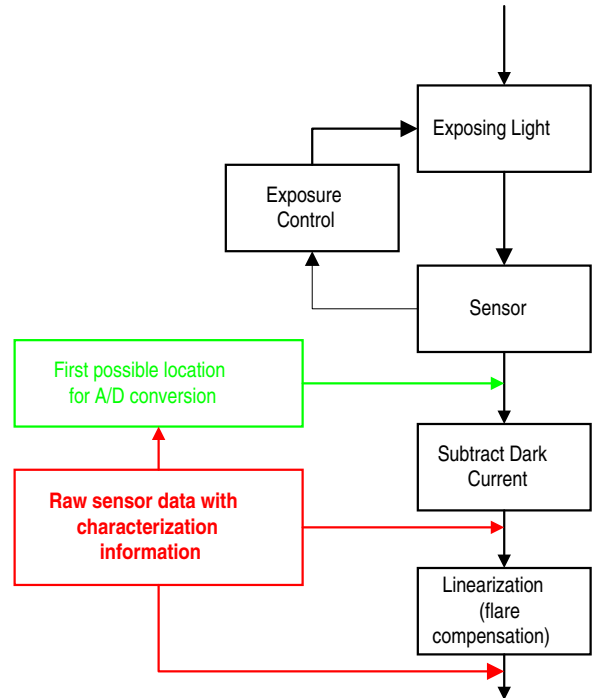


Figure 2. Encoding to sensor space

Unrendered Image Space

The transformation from sensor space to unrendered, device-independent space is image and/or device specific: linearization, pixel reconstruction (if necessary), white point selection, followed by a matrix conversion (Figure 3). If the white point of a scene is not known, as is often the case in digital photography, it has to be estimated.⁷

The purpose of an unrendered image color space is to represent an estimate of the scene’s or the original’s colorimetry. An unrendered space maintains the relative dynamic range and gamut of the scene or original.

Unrendered images will need to go through additional transforms to make them viewable or printable. Appearance modeling can be applied when an equivalent or corresponding reproduction is desired, and the output medium supports the dynamic range and gamut of the original. In most applications, the goal is to create a preferred reproduction, meaning the image is optimized to look good on a specific medium with a different dynamic range and gamut than the original. In that case, a digital photography reproduction model is applied. Unrendered image spaces can be used for archiving images when it is important that the original colorimetry is preserved so that a facsimile can be created at a later date.

The advantage of unrendered image spaces, especially if the images are encoded in higher bit-depth, is that they can always be tone and color processed for all kinds of different rendering intents and output devices at a later date. The quality of the colorimetric estimate depends on the ability to choose the correct scene adopted white point, and the correct transformations.

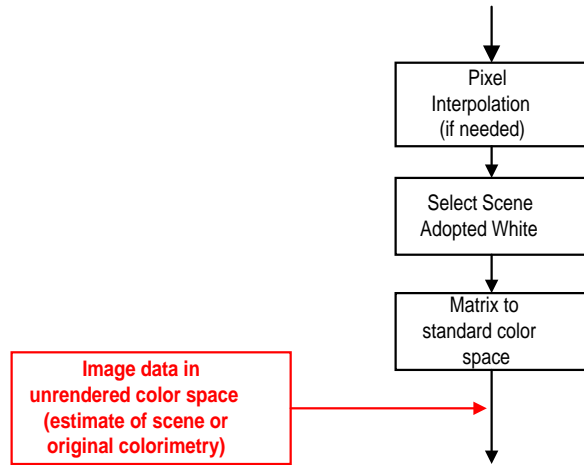


Figure 3. Encoding from sensor to unrendered space

Examples of color spaces that can describe an estimate of the scene’s or original’s colorimetry are ISO RGB, CIE XYZ, Photo YCC, and CIELAB.

Rendered Image Space

Rendered image spaces are color spaces based on the colorimetry of real or virtual output characteristics. Images can be transformed into rendered spaces from either source or unrendered image spaces. The complexity of these transforms varies: they can range from a simple video-based approach to complicated image dependent algorithms. The transforms are usually non-reversible, as some information of the original scene encoding is discarded or compressed to fit the dynamic range and gamut of the output (Figure 4). The transforms are image specific, especially if pictorial reproduction modeling is applied. The rendering intent of the image has therefore been chosen, and may not be easily reversed. For example, an image that has been pictorially rendered for preferred reproduction cannot be re-transformed into a colorimetric reproduction of the original without knowledge of the rendering transform used.^{8,9,10}

Rendered image spaces are usually designed to closely resemble some output device characteristics, ensuring that there is little loss when converting to the output specific space. Most commercial image applications only support 24-bit image encoding, making it difficult to make major tone and color corrections at that stage without incurring visual image artifacts. Some rendered RGB color spaces are designed so that no additional transform is necessary to view the images; in effect, the rendered RGB color space is the same as the monitor output space. For example, sRGB is a rendered image space that describes a real output and as such, is equivalent to an output space.

Output Space

Transforms from rendered RGB spaces to output spaces are device and media specific (Figure 5). If a rendered space is equal or close enough to real device characteristics, such as “monitor” RGBs, no additional transformation to device

specific digital values is needed. In many cases, however, there is a need for additional conversions. For most applications, this can be accomplished using the current ICC color management workflow. An “input” profile maps the reproduction description in the rendered space to the profile connection space (PCS), and the output profile maps from the PCS to the device and media specific values.

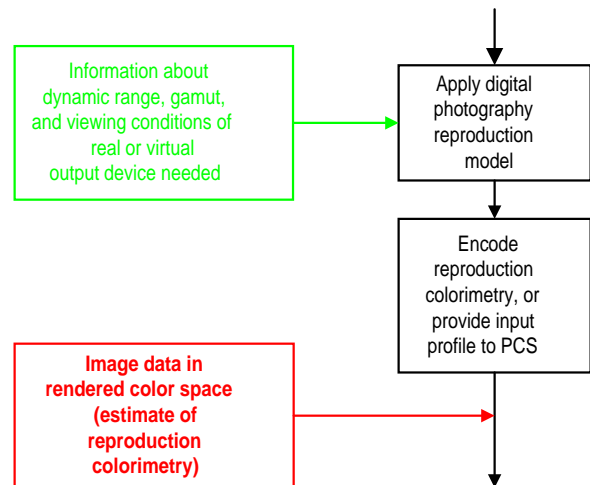


Figure 4. Encoding from unrendered to rendered space

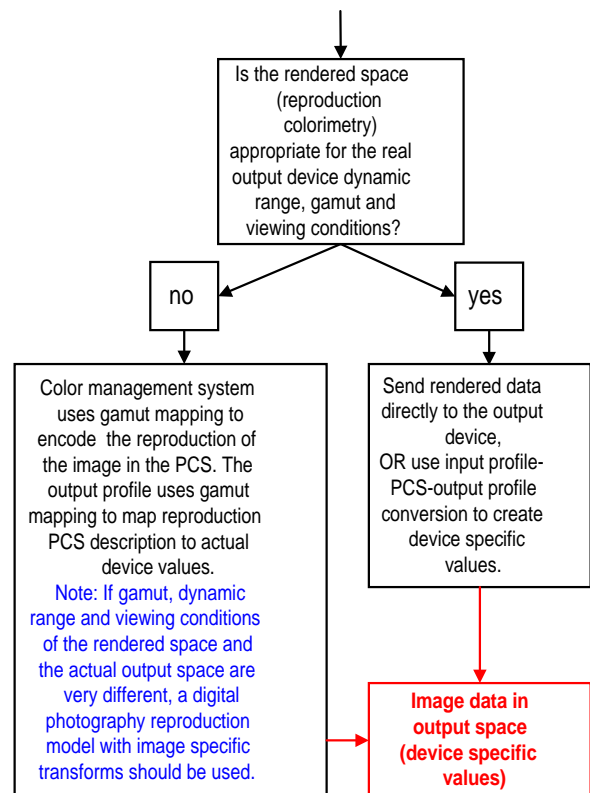


Figure 5. Encoding from rendered to output space

If the gamut, dynamic range, and viewing conditions of the rendered space are very different from those of the actual output space, it might be more advantageous to use a reproduction model that allows image specific transforms than to use the current color management system. Adjusting for different viewing conditions and dynamic range often are not implemented in current applications, and out of gamut colors require image dependent mapping for optimal reproduction.

Apart from graphic arts applications, it is rare today that images are archived and communicated in output space, such as device and media specific RGB, CMY, or CMYK spaces. However, there are many legacy files, such as CMYK separations and RGB monitor specific images that need to be color managed so that they can be viewed and printed on other devices.

Standard RGB Color Spaces

Unrendered RGB Color Spaces

There is currently only one unrendered RGB color space that is in the process of becoming a standard. ISO RGB is defined in ISO 17321 – Graphic Technology and Photography – Colour characterisation of digital still cameras using colour targets and spectral illumination.³ ISO RGB is the reference color space to evaluate digital still camera color analysis. The transformations from sensor data to ISO RGB are defined in the standard.

ISO RGB data represent an estimate of the scene or original colorimetry. There are no specified dynamic range or viewing conditions associated with ISO RGB.

When ISO RGB is encoded in 8 bits, a limited gamut is defined by the ITU-R BT.709 primaries.¹¹ Extended ISO RGB is encoded in 10-16 bits with an unlimited gamut; one bit is used as the negative sign to encode out-of-gamut colors.

Instead of using one of the traditional unrendered color spaces such as CIE XYZ, ISO RGB has been developed for the purpose of digital camera characterization because recent research indicates that white point conversions perform best when based on RGB color matching functions.¹² Linear ISO RGB values can be transformed to CIE XYZ values by a matrix conversion based on the CIE XYZ values of a reference white point. The equi-energy RGB to XYZ matrix is pre-multiplied with a diagonal matrix containing the XYZ values of the reference white point. The XYZ values associated with the ISO RGB values for the specified reference white point can then be calculated using that matrix.

Rendered Color Spaces

There are a number of rendered RGB color spaces used in different imaging applications and workflows today. Following are descriptions of the most common rendered RGB color spaces, selected based on the applications.

Multimedia: sRGB

sRGB is described by IEC 61966-2-1 as a default color space for multimedia applications.¹³ It is a rendered space, and is based on the characteristics of a CRT reference display. In its normative part, the standard defines the relationship between 8-bit sRGB values and CIE 1931 XYZ values as measured at the faceplate of the reference display. The encoding transformations do not take into account the veiling glare defined in the reference viewing conditions. The reference display white point and primaries are defined according to ITU-R BT.709.¹¹ The colorimetry seen by an observer looking at an image on a reference display in the reference viewing conditions is described in an informative annex by recommending how to encode for veiling glare.

The standard does not define how to encode image data into sRGB, how to map to a different rendered space, how to adjust for different viewing conditions, or how to map to an output space.

The purpose of sRGB is to define a rendered color space for data interchange in multimedia. Due to similarities of the defined reference display to real CRT monitors, often no additional color space conversion is needed to display the images. However, conversions are required to transform data into sRGB and then out to devices with different dynamic ranges, gamuts and viewing conditions. Because of sRGB's CRT-based gamut, which is smaller than the range of colors achievable on hardcopy reproduction, this can be problematic as few tools are available to do this properly.

Microsoft and Hewlett-Packard recently proposed sRGB64, which extends the tonal range and coding precision of sRGB.¹⁴

Editing Space: ROMM RGB

ROMM (Reference Output Medium Metric) RGB is a wide-gamut, rendered RGB color space. It was designed by Eastman Kodak and is intended as an RGB color space for manipulating and editing images after the initial rendering has been applied. The ROMM RGB primaries are not tied to any monitor specification. Rather they were selected to wholly enclose an experimentally-determined gamut of surface colors, so that there would not be any loss of color information when representing reflectance colors that had been captured in an unrendered color space. For the specific primaries chosen, a contrast-boosting tone scale mapping of the ROMM RGB images results in small or minimal hue shifts on an a*-b* plot.⁹ ROMM RGB uses a D50 white point, which is a standard for viewing and evaluating graphic arts reproductions, as well as the ICC PCS white point.

ROMM RGB has the largest gamut of the rendered RGB spaces described here. By selecting a gamut that wholly encloses most real world surface colors, many ROMM RGB values are wasted in the production of reflection hard copy, in that they do not correspond to reflectance colors and are never used. For a given number of bits, the wider the gamut, the coarser the quantization and the greater the potential for visible artifacts due to quantization and subsequent processing of the image.

Evaluations have shown that bit depth quantization at 8 bits only create visible artifacts in photographic images if the image processing is very aggressive. Still, ROMM RGB offers 12- and 16-bit encoding, in addition to the usual 8-bit encoding, to allow for greater precision. ROMM RGB can be used as a wide-gamut RGB working space in Adobe Photoshop 5.

Adobe Photoshop Working Space: Adobe RGB 98

With Photoshop 5, Adobe introduced the concept of a working space that is device independent. The goal is to make the image data more portable and not tied to the RGB monitor of anyone’s desktop. It is also the space the user will import images to from different sources and make editing decisions in.

Adobe RGB 98 was intended to provide a larger gamut so pre-press users can use it as the default working space in Photoshop 5. It is based on the SMPTE-240M standard and was later renamed Adobe RGB 98.

The potential drawback of Adobe RGB 98 space is that it includes many colors unprintable using typical CMYK printers. Savvy users can minimize this problem by picking

a target output device and limiting the color section to the ones within the output gamut.

Legacy Images: Apple RGB

Apple RGB is based on the classic Apple 13" RGB monitor. Because of its popularity and similar Trinitron-based monitors that followed, many key publishing applications, including Adobe Photoshop and Illustrator, used it as the default RGB space in the past.

Although the gamut of Apple RGB space is not much different than sRGB, this space represents many legacy files in the desktop publishing world.

Video RGB: NTSC, EBU, ITU-R BT.709

There are several RGB standards for video applications. Video standards involve more than just RGB color spaces, but the focus here will only be on color spaces. Starting with the NTSC spec in the early 50's to the latest spec for HDTV, video standards have tracked advances in the CRT technology used in broadcast television applications. The original NTSC system was geared toward a CRT display.

Table 1: Attributes of standard RGB color spaces

Color Space	Type	Encoding	Gamut	White Point	Primaries		Specified Dynamic Range and Viewing Conditions	
					x	y		
ISO RGB	Unrendered	8-bit nonlinear	Limited	floating	floating		No	
Extended ISO RGB	Unrendered	10- to 16-bit nonlinear	Unlimited (signed)	floating	floating		No	
sRGB	Rendered	8-bit nonlinear	CRT	D65	R	0.64	0.33	Yes; reference viewing environment defined, with D50 as ambient white point
					G	0.30	0.60	
					B	0.15	0.06	
ROMM RGB	Rendered	8-bit nonlinear, 12-, 16-bit optional	Wide	D50	R	0.7347	0.2653	Yes; reproduction viewing environment defined
					G	0.1596	0.8404	
					B	0.0366	0.0001	
Adobe RGB 98	Rendered	8-bit nonlinear	Extended CRT	D65	R	0.64	0.34	No
					G	0.21	0.71	
					B	0.15	0.06	
Apple RGB	Rendered	8-bit nonlinear	CRT	D65	R	0.625	0.34	No
					G	0.28	0.595	
					B	0.155	0.070	
NTSC RGB	Rendered	Nonlinear	CRT	Ill. C	R	0.67	0.33	partial gamma correction to compensate for destination viewing conditions
					G	0.21	0.71	
					B	0.14	0.08	
EBU RGB (CCIR 601)	Rendered	Nonlinear	CRT	D65	R	0.64	0.33	No
					G	0.29	0.60	
					B	0.15	0.06	
ITU-R BT.709	Rendered	Nonlinear	CRT	D65	R	0.64	0.33	No
					G	0.30	0.60	
					B	0.15	0.06	

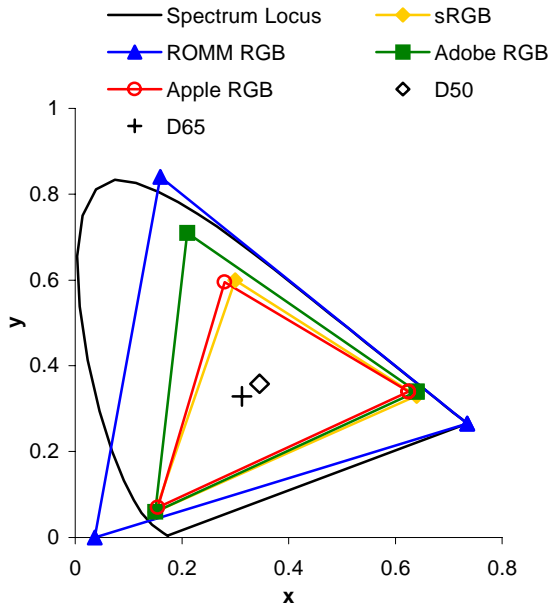


Figure 6. x,y -Chromaticity diagrams for sRGB, ROMM RGB, Adobe RGB 98 and Apple RGB

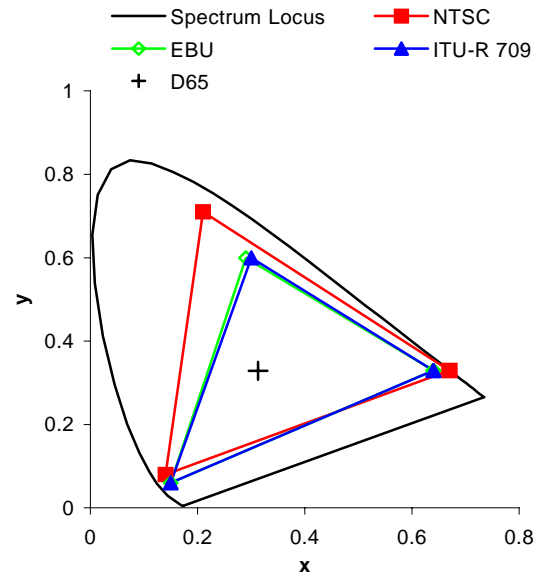


Figure 7. x,y -Chromaticity diagrams for video RGB (NTSC, EBU, and ITU-R BT.709).

The PAL system for color television came along after NTSC. The EBU decided to standardize the PAL system on primaries based on the displays available at the time. The CCIR (now the ITU-R) made this a standard in 1974.¹⁵

The latest video RGB standard is ITU-R BT.709 for the production and exchange of HDTV programming.¹¹

The 601 and 709 phosphors are practically the same, but the 709 standard defines a different form of gamma correction (transfer characteristic). Compared to NTSC displays, modern CRTs have brighter and more efficient phosphors.

These benefits have come at a slight cost, since the NTSC display could produce purer yellows and reds. However, the 709 primaries are preferred, as they match most modern CRTs. sRGB is based on the 709 primaries.

RGB Color Space Implementation Issues

The different standard color spaces discussed above were all developed for specific applications. Which RGB space to choose when the application does not fall into a well-defined imaging workflow can be difficult. Following are a few points that need to be considered.

Sensor, Unrendered, Rendered, or Output Color Spaces

For high-end archiving purposes, storing images as raw sensor data with the necessary sensor and illumination characteristics is preferred. Any further transformations are dependent on current engineering practices and knowledge, which might be improved in the future. However, it is rare to find independent imaging applications that can read any raw image data and convert them using the necessary sensor and scene characteristics.

In most cases, the “input” software and hardware for scanners and digital cameras integrate the transformation to rendered color space without allowing access to any intermediate data form. Most scanning software allows some manual intervention in the process, therefore allowing a user to manually manipulate the reproduction model.

If raw sensor data can be saved, the file format is either proprietary, allowing further processing only with the manufacturer’s own software, or the necessary sensor and scene information is not included. There is some scanning software that allows the user to save the image in an unrendered color space such as CIELAB or Photo YCC. Archiving in an unrendered color space is feasible. However, the quality of the colorimetric estimate of the scene is dependent on the scene content, capture device, and transformation.

In most applications today, images are converted into a rendered space for archiving and data interchange. That practice allows the owner/creator of an image to define the initial rendering, but these transformations are frequently not easily reversible, because the reverse transform and tools to implement it are not available. Choosing the right rendered space, depending on the future use of these images, is critical.

Gamut Size

If editing and output specifications are not known, the rendered color space for archiving and data interchange should have a wide gamut. However, wide gamut color spaces cannot be viewed without additional conversions on today’s monitor, and these conversions may need to be image specific. Many image-viewing applications do not enable automatic conversions, either making wide gamut

RGB images look flat and de-saturated, or clipping out of gamut colors. If no color space conversion is feasible before viewing, the interchange space's gamut should be as close as possible to the output space.

Encoding

Linear encoding (in intensity) is acceptable when high-bit depth information can be retained, and file size doesn't matter. In most cases, a nonlinear, perceptually compact encoding (nonlinear in intensity, but linear in lightness or brightness) is preferable, since the visual artifacts due to image processing would be equally visible across the tone scale.

A system that uses a linear sensor for capture and a CRT for output needs gamma correction somewhere in the imaging chain in order to produce an overall linear system. Gamma correction compensates for the non-linear power-law transfer characteristic of the CRT. As it turns out, a gamma corrected signal looks much like a perceptually compact encoding and offers the same benefits. Gamma corrected signals are input directly to the monitor, without further processing, to give the desired output. Partially-gamma corrected signals are used to obtain a specific visual effect. For example, the NTSC system only does a partial gamma correction so that the overall system's transfer characteristic has a contrast greater than unity to produce a more preferred appearance and reproduce actual appearance in dim viewing conditions. The discussion of gamma merits an entire paper in itself; see for example.¹⁶ Gamma is an important implementation issue for RGB color spaces.

If a color space has a wide or unlimited gamut, 8-bit encoding might not be enough. Banding effects can appear, depending on image color distribution, editing, and/or color space conversion. However, 16 bit/component RGB is not widely supported yet in either applications or file formats. Images that will go through extensive image processing and color space conversions should be encoded in higher bit-depth. If that is not possible, the size of the gamut should be reduced.

Color Space Conversions

Converting in and out of different color spaces can cause severe image artifacts. The more mismatched the gamuts and white points are, the stronger the effects. Most current implementations of converting to RGB spaces clip out-of-gamut colors instead of mapping them more intelligently. If different color spaces used in a workflow have different white points, either the image application has to perform a white point conversion, or it has already been built into the profile representing the color space. These operations can also contribute to artifacts. Other factors such as adjusting for different viewing conditions and dynamic ranges are often unspecified by the color space definition or unimplemented by imaging applications.

Compression

The first commercial television system was the CBS Sequential Field System, which, for a short time in the early

50's, was the US color television standard. This system was built around a sequential color display, which without memory or signal processing, required the transmission of RGB signals. One of the reasons that this system failed and was shortly replaced by the NTSC system was that it was wasteful of bandwidth: a better quality image could be obtained with the same bandwidth by transmitting luminance-chrominance instead of RGB values.¹⁷ Two reasons why RGB is disadvantaged for compression for image interchange is that RGB image separations are highly correlated and unable to take significant advantage of the spatial resolution reduction techniques that are applied in luminance-chrominance systems. Hence, color television, color facsimile and the Web use of JFIF—all of which emphasize efficient image interchange—use luminance-chrominance-type signals, instead of RGB.

Supporting Applications and Formats

Applications and file formats alike, there are two kinds of support for standard color spaces: one is the built-in kind where the RGB spaces are supported by name. No profile or additional information is necessary. The other kind is by specification, where typically a profile or a set of parameters describe the color space.

Color management systems such as ColorSync, ICM, or Postscript support RGB spaces typically by specification. Color values can be converted in and out of RGB spaces as long as an accurate profile is provided for that space. Some file formats also support RGB spaces by specification. PDF, TIFF, JPEG, PICT, EPS, and PNG all have ways to associate profiles to further define the RGB color spaces. Few formats have built-in support for the standard color spaces that are discussed here. FlashPix, MNG (Multi-image Network Graphics) and HTML 4.0 support sRGB—MNG Ver. 0.96 allows sRGB “chunks,” and RGB color data in HTML 4.0 is defined in sRGB color space, which is effectively the default color space. Applications such as Photoshop support some standard color spaces by name, such as Adobe RGB, Apple RGB, sRGB, etc., as well as by specification in the form of a profile.

The Future of RGB Display Spaces

Many standard RGB color spaces, such as ITU-R BT.709 RGB or sRGB, assumed a standard monitor, based on phosphor/monitor characteristics at the time, and applications were built to deliver RGB matched to that monitor. It was a world without variability, or if you like, with uniformity. It was easier to enforce then than it is now. Some of the conversions from that era have persisted, even though monitors and phosphors have changed.

With new display technologies maturing, this CRT-centric view will increasingly find that the underlying assumptions are no longer true. Using TFT-based flat panel as an example, there are key differences to consider.

No Standard Phosphor Sets

In the CRT based displays, there are about six phosphor sets that are used, with P22 and EBU being the most popular ones. Their primaries are fairly close to each other. However, flat panel displays do not have standard sets of filters and the difference between primaries across different devices can be quite large.

Gamma Curves

Most RGB spaces define gamma using a typical CRT based model of gamma-offset-gain or even pure power functions. However, most flat panel displays have an S-shaped gamma curve. Fitting it into the current model means significant errors in either the highlight or shadows. The alternative is to distribute the errors more evenly. In either case, it does not offer good color accuracy.

Viewing Conditions and Dynamic Range

Many RGB spaces define viewing conditions and dynamic range based on the limitations of CRT monitors. Because of the limited dynamic range and brightness level of CRTs, the typical viewing condition is a dim surround. However, the latest flat panel displays can output 230+ cd/m² with a dynamic range over 300 to 1. The contamination of ambient light is also minimized. The viewing conditions in future work environment could be very different from those today.

Considering these points, clearly the approach of defining a standard RGB space that is similar to CRT monitor characteristics to avoid output device specific conversions will not always be feasible in the future. With the advances in flat panel displays and other display technologies, the future is fast approaching.

Conclusions

There is no "one size fits all" approach, no one RGB color space that is ideal for the archiving, communicating, compressing, and viewing of color images. The correct color space, be it RGB or not, depends on the application. For archiving, the first consideration should be the future use of the image. If a single, final use of the image is known, then a rendered image space closely related to the intended use may be selected. On the other hand, if the desired rendering intent is known, but more than one type of output is desired, then use of a wide-gamut rendered image space may be the best choice. If the greatest flexibility is desired, then a high bit depth sensor or unrendered image space should be selected.

Acknowledgments

We wish to thank Ann McCarthy of Xerox Architecture Center and Jack Holm of Hewlett-Packard Laboratories for their careful review and valuable inputs.

References

1. J. King, "On the Desktop Color System implied by the ICC Standard." *IS&T/SID Proc. 6th Color Imaging Conference*, pp. 249-252, 1998.
2. CIE Publication N. 17.4 – 1987, *International Lighting Vocabulary*, Definition 845-03-25.
3. ISO 17321, WD 4, Graphic Technology and Photography – *Colour characterisation of digital still cameras (DSCs) using colour targets and spectral illumination*, November 1999. At www.pima.net/it10a.htm
4. S. Süssstrunk and J. Holm, "Camera Data Sheet for Pictorial Electronic Still Cameras." *SPIE Vol. 2416*, pp. 5-16, 1995.
5. J. Holm, "Issues Relating to the Transformation of Sensor Data into Standard Color Spaces," *IS&T/SID Proc. 5th Color Imaging Conference*, pp. 290-295, 1997.
6. ISO 12234-2 DIS Photography - *Electronic still picture cameras – Removable memory - Part 2: Image data format - TIFF/EP*. At www.pima.net/it10a.htm
7. P. M. Hubel, J. Holm, G. D. Finlayson, "Illuminant Estimation and Color Correction," *Proc. Colour Imaging in Multimedia – CIM98*, Derby, England, pp. 97-105, 1998.
8. J. Holm, "Photographic Tone and Colour Reproduction Goals." *Proc. CIE Expert Symposium '96*: pp. 51-56, 1996.
9. Eastman Kodak Company, "Reference output Medium Metric RGB Color Space (ROMM RGB) White Paper," Version 2.2, 1 July 1999.
10. S. Süssstrunk, "Imaging Systems for the Creation of a Large Digital Image Archive," *Proc. IS&T's 50th Annual Conference*, pp. 322-325, 1997.
11. ITU-R Recommendation BT.709-3: 1998, *Parameter values for the HDTV standards for production and international programme exchange*.
12. G. D. Finlayson and M. S. Drew, "White-point preserving color correction," *Proc. IS&T/SID 5th Color Imaging Conference*, pp. 258-261, 1997.
13. IEC 61966-2-1/FDIS: Multimedia systems and equipment – *Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB*.
14. www.sRGB.com/sRGB64
15. W. N. Sproson, *Colour Science in Television and Display Systems*, 1983, Adam Hilger, Bristol.
16. C. Poynton, "The rehabilitation of gamma, in Human Vision and Electronic Imaging III," *SPIE*, vol. 3299, pp. 232-249, Jan. 1998; see edited version at www.poynton.com/~poynton
17. W. F. Schreiber, R. R. Buckley, "Introduction to Color Television - Part 1," *IEEE*, vol. 87, no.1, pp. 173-179, January 1999.