Requirements for Unambiguous Specification of a Color Encoding: ISO 22028-1

Kevin Spaulding
Eastman Kodak Company
Rochester, New York

Abstract

In recent years, many different color encodings have been developed for use in a wide variety of digital imaging applications. Some of these color encodings have been proposed/standardized by various standards bodies or industry groups, while others are being used without any official standardization. In order to unambiguously communicate digital color images in open systems, it is important that the color encoding used to represent the image data be clearly and completely defined. In many cases, the lack of such clear specifications has led to confusion and inconsistency in the way some color encodings have been used. This can lead to significant interoperability problems when various components in an imaging system interpret the color encoding in an incompatible manner. Recently, ISO 22028-1 (a joint project between ISO TC42, ISO TC130 and CIE) has reached the Committee Draft stage. The goal of this standard is to define a framework and a consistent set of guidelines that can be used for defining digital color encodings. While this standard does not itself define any color encodings, it is anticipated that it will be used as the basis for specifying future color encoding standards.

Introduction

Modern digital imaging systems serve a variety of consumer and commercial applications. A fundamental choice for any imaging system architecture is how to represent images numerically: in what color space and with what digital encoding. Depending on the application, differing priorities will apply to such system attributes as image quality, interoperability, simplicity of system architecture, efficiency of image computations, and the flexibility for optimally using images for a variety of purposes. Trade-offs among these attributes will be application-dependent.

In some applications, a single-color encoding designed to be compatible with the prevalent mode of image viewing by the end-user may suffice. Because both multimedia and Internet-based imaging rely heavily on the viewing of images on a softcopy display, the use of the *sRGB*[1] color encoding makes sense for many of those applications, because it is based on the response of a standard CRT. However, because the color gamut of *sRGB* does not encompass the color gamuts of many common input and output devices, a system architecture that depends exclusively on the use of *sRGB* would represent an unacceptable compromise in color reproduction accuracy for some applications. As a result, there is a need for a flexible framework that can support the definition of a variety of color encodings (extended-gamut color encodings, as well as display-referred color encodings like *sRGB*) required to meet the needs of the full range of common applications and workflows. Because many digital imaging workflows include image representations that do not directly correspond to the color of an output picture (e.g., they may be encodings of photographic negatives or original scenes that still need to be color rendered in order to determine appropriate picture colors), it is important that this framework be able to support the definition of color encodings appropriate for a variety of image states,[2,3] and for all stages of the workflow, from the initial image capture through to the final step of producing a softcopy or hardcopy reproduction. These needs are addressed in ISO 22028-1[4] (Photography and graphic technology—Extended colour encodings for digital image storage, manipulation and interchange—Part 1: Architecture and requirements), which defines the requirements for unambiguously specifying a color encoding, together with a generic image state architecture that provides a framework for the classification and use of various color encodings.

Color Encoding Hierarchy

Terms like “color space” and “color encoding” have been used rather loosely, and often inconsistently, by different individuals, and even in various color encoding standards. During the course of developing ISO 22028-1, it was found that this inconsistency in terminology presented a significant barrier to converging on a consensus solution. What was sufficient for defining a color encoding according to one usage of the term, was insufficient according to another definition. Therefore, one of the first steps was defining a consistent set of terms and definitions. The result was a hierarchy of terms building from a color space, up to a color image encoding as is shown in Fig. 1.

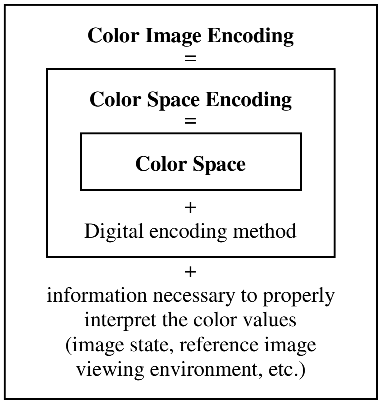

Figure 1. Hierarchical relationship between color space, color space encoding, and color image encoding.

A *color space* is simply defined to be a "geometric representation of colors in space, usually of three dimension," which is right out of the official CIE vocabulary.[5] Examples of such color spaces would be *CIE XYZ*, *CIELAB* or continuous (unquantized) *RGB* values for a particular set of primaries. Defining a color space doesn't say anything about how colors are represented digitally, nor does it say anything about how to interpret the color values. However, it does include the specification of any nonlinear transfer functions that are to be applied to the color values, as this will change the "geometric representation" of the color values. For example, a radiometrically linear *CIE XYZ* color space would be a different color space than a logarithmic *CIE XYZ* color space.

For digital imaging applications, it is also necessary to define a digital *color encoding* of the color space. This is where the greatest confusion and inconsistency in terminology existed. To resolve these inconsistencies, two different qualifiers were added to define two distinct levels in the hierarchy. A *color space encoding* was defined to be a "digital encoding of a color space, including the specification of a digital encoding method, and an encoding range." This is simply a specification of how the color space values are represented digitally. This could include the definition of an integer representation of a specific bit-depth, or even a digital floating point representation.

Finally, a *color image encoding* is defined to be a "digital encoding of the color values for a digital image, including the specification of a color space encoding, together with any information necessary to properly interpret the color values such as the image state, the intended image viewing environment, and the reference medium." In order to clearly and unambiguously communicate the color of digital images in an open system, it is important that the color be specified by a complete color image encoding, rather than simply defining a color space, or even a color space encoding. Accordingly, it is also important that any color encoding standards that are defined should specify complete color image encodings. The main body of the ISO 22028-1 proposal is devoted to specifying what information should be included in a color image encoding definition in order to provide an unambiguous color representation, as well as providing a generic architecture within which various color space encodings can be described.

Image State Architecture

The need for various color encodings and the rationale for their specifications can be best understood in the context of the particular industry and workflow for which they are intended. While the details of the particular workflows are very diverse and often complex, a simple model can be used to represent their core activities as shown in Fig. 2. This high-level image processing architecture, and the associated terminology, facilitates a common framework for classifying different color encodings, and describing imaging chains for most common types of digital imaging systems. Correctly understanding where a particular color image encoding fits within this framework is important for determining the appropriate interpretation and use of the image data.

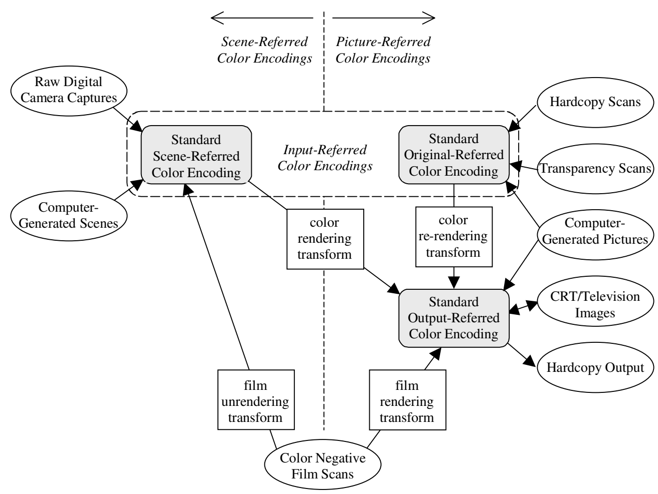

Figure 2. Image state diagram showing relationship between various types of color encodings.

It should be noted that while this image processing architecture shows examples of where different types of devices might fit within typical workflows utilizing standard color encodings, this standard does not define any workflows that must be used for any particular application, nor should it be considered to constrain the workflows to those shown in this figure. For example, raw digital camera captures may be processed directly to an output-referred image state without passing through a standard scene-referred color encoding. It should also be noted that a particular workflow might include additional color encodings corresponding to image states different from the standard image states defined in this architecture. For example, it may be useful to define a standard color encoding for representing color negative scans, or an intermediate color encoding for partially color rendered images. *While color encodings in these non-standard image states may be valuable internal to particular applications, they should generally not be used for image interchange in an open systems environment unless there is consensus on the definition of a corresponding image state*.

The image state diagram in Fig. 1 shows that color encodings can broadly be categorized as either *scene-referred* or *picture-referred*. *Scene-referred color encodings* are representations of the estimated color space coordinates of the elements of an original scene, where a scene is defined to be the spectral radiances of a view of the natural world as measured from a specified vantage point in space and at a specified time. On the other hand, *picture-referred color encodings* are representations of the color space coordinates of a hardcopy or softcopy image.

Scene-referred image data may be represented by an encoding of scene colorimetry using a device-independent color space such as *CIE XYZ* or *CIELAB*, or may be represented in terms of the response of an idealized scene capture device such as *RIMM RGB*,[6] or *scRGB*.[7] Scene-referred image data may correspond to an actual view of the natural world, or a simulation of such a view. It may also correspond to a modified scene determined by applying modifications to an original scene. For example, such modifications could include removing haze from the captured image, or allowing a user to manually adjust the exposure/white balance. It could also include more complex operations, such as using a “dodge-and-burn” algorithm, to correct overexposed regions of a backlit scene. (This can be viewed as being analogous to “re-lighting” the scene.) Scene modifications could also include applying desired changes to the scene such as simulating a “night” scene, making grass greener to make it look healthier, or making the sky bluer to make it look clearer. Any such scene modifications should leave the image in a scene-referred image state, and should be done in the context of the expected color rendering transform. For example, typical color rendering transforms will include a boost in the chroma of the image. Any boost in colorfulness of the scene (e.g., making the grass greener) should be done with the knowledge that there will be an additional chroma boost during color rendering. Consequently, the color rendering transform should be included in any image preview path that is used to provide subjective feedback to a user during the scene-editing process.

It should be noted that the image colorimetry of the scene-referred image may contain inaccuracies due to factors such as dynamic range limitations of the capture device, noise from various sources, quantization, optical blurring and flare, and color analysis errors caused by capture device metamerism. In some cases, these sources of inaccuracy can be significant.

Picture-referred color encodings can be further subdivided into original-referred color encodings and output-referred color encodings. *Original-referred color encodings* are representative of the color space coordinates (or an approximation thereof) of a two-dimensional hardcopy or softcopy image. Original-referred image data is generally produced by scanning artwork, photographic transparencies/prints, or photomechanical reproductions, etc. The characteristics of original-referred image data are tightly coupled to the characteristics of the original image source. Original-referred image data may be represented by an encoding of picture colorimetry, in terms of the response of an idealized measurement device such as a *Status A* densitometer, or in terms of device-dependent control signals for a particular image capture device.

Output-referred color encodings are representative of the color space coordinates of image data that is appropriate for a specified real or virtual output device and viewing conditions. Standard output-referred color encodings are color encodings that are referred to a particular standardized real or virtual output device and viewing conditions. Image data intended for open interchange is most commonly standard output-referred. Output-referred image data may be represented by an encoding of picture colorimetry, by a color encoding derived from CIE colorimetry such as *ROMM RGB*[8] or *bg-sRGB*,[9,10] by the control signals for a standardized real or virtual output device such as *sRGB*[1] or *SWOP CMYK*,[11] or by device-dependent control signals for a particular softcopy or hardcopy output device.

Note that output-referred image data may become the starting point for a subsequent reproduction process. For example, *sRGB* standard output-referred image data are frequently considered the starting point for many desktop printing systems. In this case, the printing systems will generally perform a color re-rendering process to transform the *sRGB* color values to those appropriate for the particular output device and assumed viewing conditions.

A *color rendering transform* is used to transform a scene-referred image to a corresponding output-referred image. Color rendering transforms embody the tone and color reproduction aims of an imaging system, relating the corresponding scene color values to the desired picture color values. Color rendering transforms are usually substantially different from identity transforms and accomplish several important purposes including compensating for differences in the input and output viewing conditions, applying tone and gamut mapping to account for the dynamic range and color gamut of the output medium, and applying color adjustments to account for preferences of human observers.[2] Color rendering transforms are typically proprietary and irreversible.

For cases where original-referred image data is to be transformed to output-referred image data, a *color re-rendering transform* should generally be used to adjust the image colorimetry when the characteristics of the media and/or intended viewing condition are not the same. For example, if an original-referred image represents the colorimetry of a photographic transparency intended to be projected in a dark room and it is desired to transform the image data to an output-referred color encoding associated with a reference reflection print medium and viewing conditions, then the color re-rendering transform would need to account for the difference in the intended viewing conditions, as well as the differences in the media dynamic range and color gamut.

Both scene-referred color encodings and original-referred color encodings can be grouped together into a category referred to as *input-referred color encodings*. Although scene-referred image data and original-referred image data are both subsets of the input-referred image data classification, the two types of image data have significantly different characteristics, and should generally be treated quite differently. Therefore, it usually will not be appropriate to use the same color image encoding for both types of data. As a result, standard color encodings for input-referred image data should be designated to be scene-referred or original-referred rather than simply input-referred.

Color negative film scans can be transformed to either a scene-referred image state or an output-referred image state. Because of the extremely wide dynamic range of film, it represents an excellent means for capturing scene-referred image data. To back out the characteristics of the film, a *film unrendering transform* is needed to infer the original scene colors from the film scan. Alternatively, the film scan can be used to directly form an output-referred image. In this case, a *film rendering transform* is needed to determine the desired picture colorimetry from the film scan. Typically, such transforms are designed to model the response of a real or idealized photographic paper exposed using a conventional photographic enlarger.

Information Needed to Define a *Color Space*

The information needed to define a color space will depend on the type of color space being specified. Color spaces can be broadly categorized into three types: colorimetric, color appearance, and device-dependent.

Colorimetric Color Spaces

For colorimetric color spaces, it is necessary to define the relationship between the color values and CIE colorimetry. In some cases, the color space may be one of the standard CIE color spaces such as *CIE XYZ* or *CIELAB*. In those cases, the color spaces are themselves measures of CIE colorimetry. In other cases, a color space may be specified that has a defined relationship to CIE colorimetry. In those cases, the transformation back and forth between the color space and standard CIE colorimetry should be defined. Because colorimetric quantities will be dependent on measurement conditions, the assumed measurement conditions associated with all CIE colorimetry should be specified. (e.g., 0/45 geometry vs integrating sphere, etc.) The following is an outline of the information needed to define several common types of colorimetric color spaces

- CIE color spaces
  - Color space (e.g., *CIE XYZ*, *CIELAB*, etc.)
  - Color space white point
- Additive RGB color spaces
  - Chromaticity values for RGB primaries
  - Color space white point
  - Color component transfer function (e.g., gamma function, etc.)
- Luma-chroma color spaces derived from additive RGB color spaces (e.g., YC_RC_B)
  - Additive RGB color space
  - Luma-chroma matrix

Color Appearance Color Spaces

Color spaces based on color appearance models, such as CIECAM97s, can also be useful for encoding images in certain applications. The definition of a color appearance color space should include the specification of the particular color appearance model upon which the color space is based, together with the particular color coordinates to be used to represent the image data. (Color appearance color spaces are generally derived from CIE colorimetry, together with additional parameters relating to the image viewing environment.)

Device-Dependent Color Spaces

Some color spaces do not have a direct relationship to CIE colorimetry, but rather are defined by the characteristics of a real or idealized imaging device. There are two main classes of device-dependent color spaces. *Input-device-dependent color spaces* are defined by specifying the characteristics of a reference image capture device, including the spectral sensitivity and color component transfer function. *Output-device-dependent color spaces* are defined by specifying the relationship between the control signals of a reference output device and the corresponding output color values.

Examples of device-dependent color spaces would include densitometric color spaces such as standard Status A densities, and output-device color spaces such as printer CMY(K). Device-dependent color spaces can be defined to correspond to an actual imaging device, or to some idealized imaging device. (Output-device-dependent color spaces defined by a simple functional relationship to CIE colorimetry, e.g., *sRGB*, should be categorized as colorimetric color spaces.)

Information Needed to Define a *Color Space Encoding*

Defining a color space encoding requires the identification of a color space as well as a digital encoding method specifying the relationship between the continuous color space values and the digital representation of those values. Digital encoding methods will typically be integer digital encoding methods, although floating-point digital encoding methods will be useful in certain applications. The definition of an integer digital encoding method shall include the specification of an color space value range and a corresponding digital code value range. (Note that the digital code values should be linearly related to the continuous color space values. If a nonlinear encoding of the color space values is desired, the appropriate nonlinearity should be incorporated into the definition of the color space so that a linear quantization function can be used.)

Information Needed to Define a *Color Image Encoding*

Defining a color image encoding requires the identification of a color space encoding, together with any information necessary to properly interpret the color values such as the image state, the reference image viewing environment, and the set of valid color space encoding color values. Color image encodings intended for an output-referred image state should also define the characteristics of a reference imaging medium.

Multiple color image encodings can be associated with a single color space encoding, where the different color image encodings may differ in attributes such as image state and/or reference image viewing environment. For example, a scene-referred *CIE XYZ* color image encoding can be defined where the *CIE XYZ* values correspond to the colors of an original scene, and an output-referred *CIE XYZ* color image encoding can be defined based on the same color space encoding where the *CIE XYZ* values correspond to the colors of an output picture. The distinction in image state is important for the proper interpretation of the color values because images stored in these color image encodings should generally be treated differently. In the first case, a color rendering function would need to be applied to determine the color values for a corresponding output picture, whereas in the second case the color values would already represent the colors of an output picture.

In some cases, information such as the reference image viewing environment and reference imaging medium characteristics can be defined for a particular color image encoding so that it is fixed for all images to be stored in that color image encoding. In other cases, it may be desirable to allow this information to vary on an image-by-image basis. In such cases, metadata should be associated with the image data specifying the appropriate information.

The following is an outline of the information needed to define a color image encoding:

- Color space encoding (see previous section)
- Image state (e.g., scene-referred, output-referred, etc.)
- Reference image viewing environment
  - Image surround
  - Assumed adapted white point
  - Luminance of adapting field
  - Viewing flare
  - Set of valid color values
  - Reference imaging medium (for output-referred color image encodings)
    - Reference imaging medium white point
    - Reference imaging medium black point
    - Rendering target color gamut

Other Information in ISO 22028-1

In addition to providing the framework for the unambiguous specification of a color encoding, ISO 22028-1 also includes a number of informative annexes that are useful in understanding how this framework can be applied to real imaging systems. Annex A looks at a series of different workflows that are commonly encountered in digital imaging systems, and shows how these workflows can be mapped onto the image state architecture described in the standard. Examples of these workflows include consumer and professional digital photography, copy stand photography, scanning hardcopy, scanning color negatives, and video imaging systems. For illustration purposes, an example workflow appropriate for high-end digital photography applications is shown in Fig. 3. This workflow involves the transformation of the raw digital camera captures to a standard scene-referred color image encoding. At that point, scene-editing operations may be applied to form a "preferred scene." A color rendering transform is then applied to determine a corresponding output image in terms of a standard output-referred color image encoding. Additional picture editing may be applied at that point before sending the image to an output display.

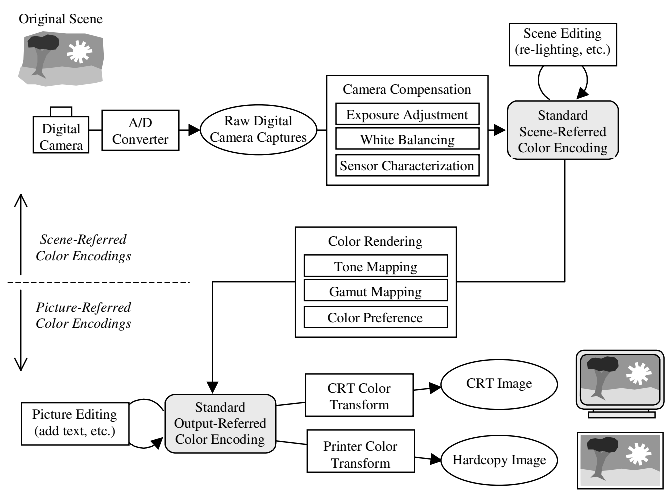

Figure 3. Generic workflow for digital photography.

Annex B compares the characteristics of a number of color encodings that have been defined by various standards organizations in terms of the requirements set forth in this standard. The color encodings that are compared include a family of color encodings based on the Rec. 709 RGB primaries (*sRGB*, *sYCC*, *bg-sRGB*, *scRGB*, *scRGB-nl*, and *scYCC-nl*), a family of color encodings based an a set of wide gamut primaries (*ROMM RGB*, *RIMM RGB*, *ERIMM RGB*), as well as two variations of the International Color Consortium's Profile Connection Space (*ICC PCS*).

The best color image encoding(s) to select for use in any particular workflow will depend many factors. Annex C discusses a number of different criteria that can be used to evaluate the alternatives to determine the optimal color image encoding for a particular application. These criteria relate to attributes such as color gamut, quantization error, perceptual uniformity, compressibility, and the complexity of transforms to other important color encodings. The relative importance of each of these criteria may vary on an application-by-application basis. Although ISO 22028-1 only summarizes these criteria at a very high level, CIE TC8-05 is currently in the process of developing detailed metrics for many of these criteria.[12]

Conclusions

This paper has summarized the work that is underway in ISO TC42/TC130 to develop a new standard defining the requirements for the unambiguous definition of a color encoding for digital images.[4] This standard defines a framework and a consistent set of guidelines that can be used for defining digital color image encodings. One aspect of this standard is the definition of a generic image state architecture that can be used to classify different color encodings. While this standard does not itself define any color encodings, it is anticipated that it will be used as the basis for specifying future color encoding standards.

Acknowledgements

The author would like to acknowledge the contributions of the many individuals who have participated in the development of ISO 22028-1, both as official committee members as well as informal reviewers. While there are far too many individuals to mention by name, deserving special note are Bill Donovan and Sabine Susstrunk who were the original project leaders for this standard, as well as Po-Chieh Hung and Jack Holm who are currently serving with the author as co-project leaders.

References

1. "Multimedia Systems and Equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB," IEC 61966-2-1 (1999).
2. J. Holm, "Issues relating to the transformation of sensor data into standard color spaces," in *Fifth Color Imaging Conference*, 1-6 (1997).
3. K. E. Spaulding, G. J. Woolfe and E. J. Giorgianni, "Optimized extended gamut color encodings for scene-referred and output-referred image states," *J. Imaging Sci. Technol.* **45**, 418-426 (2001).
4. "Photography and graphic technology—Extended colour encodings for digital image storage, manipulation and interchange—Part 1: Architecture and requirements," ISO/CD 22028-1 (2002).
5. CIE Publication 17.4, 845-03-25
6. "Photography – Electronic still picture imaging – Reference Input Medium Metric RGB color encoding: RIMM-RGB," I3A 7466:2002.
7. "Multimedia Systems and Equipment – Colour measurement and management – Part 2-2: Colour management – Extended RGB colour space – scRGB," IEC 61966-2-2, draft FDIS (2002).
8. "Photography – Electronic still picture imaging – Reference Output Medium Metric RGB color encoding: ROMM-RGB," PIMA 7666:2001.
9. "Amendment 1 to Multimedia systems and equipment - colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB," IEC 61966-2-1, draft FDIS (2002).
10. "Photography – Electronic still picture imaging – Extended sRGB color encoding – e-sRGB," PIMA 7667:2001.
11. Graphic technology—Color characterization data for type 1 printing, ANSI/CGATS TR 001 (1995).
12. "CIE TC8-05 color encoding criteria, draft 0.6," http://www.colour.org/tc8-05/documents.html (2001).

Biography

Kevin Spaulding received a BS in Imaging Science from Rochester Institute of Technology in 1983, and MS and PhD degrees in Optical Engineering from the University of Rochester in 1988 and 1992, respectively. He has been with Eastman Kodak Company since 1983 where he is currently a Senior Principal Scientist in the Imaging Science Division. He serves on several international standards committees that are working on the unambiguous communication of color images. He is project leader for the ISO 22028-1 standard, and is technical secretary for the CIE TC8-05 committee. His research interests include color reproduction, digital halftoning, image processing algorithms for digital camera and printers, and image quality metrics.