Color Encodings for Image Databases

Sabine Süsstrunk, Audiovisual Communications Laboratory,
Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland

ABSTRACT

There are a number of color encoding standards that have lately been developed and are being proposed to the imaging community for different encoding steps in the color imaging workflow, optimized for different purposes. In this paper, their applicability for master and derivative image file encoding for image databases is discussed. Several types of image databases are defined, and the advantages and disadvantages of archiving master files after different image workflow steps are outlined. Several input- and output-referred encodings are summarized and discussed in relation to their applicability for image database file encoding.

Keywords: Image Databases, Master and Derivative Image Files, Color Encodings

1. INTRODUCTION

Image database is a very general term for describing a repository of digital images, deployable for various purposes. An image database may contain any number of images, from a few hundred to millions of image objects. These image files can either *represent* digital reproductions of original artwork, such as photographs or paintings, or can *be* the original digital artwork, such as if an original scene was captured with a digital camera. It is usually implied that images should be accessible and usable over a long time frame, and should keep their value to the user indefinitely. The database should also be easy to manage, i.e. digital images should be deployed, processed, converted, and transmitted automatically. All this implies that certain strategies of initial image capture and processing techniques and a certain standardization of the workflow should be well defined before an image database project begins.

The *usage* of the digital image files ultimately dictates image quality and the necessary processing steps to be taken before archiving and delivering images. The higher the quality, the more expertise, time, and cost are associated to the generation, processing and storage of image files. While this paper focuses primarily on color encoding issues relating to the capture, maintenance, and deployment of digital images in a database, the same applies, of course, to other image database issues, such as capture and maintenance of text and administrative metadata, storage issues like file format and media, and other image processing steps like noise filtering and sharpening.

Following are three cases of most common digital image usage that define quality criteria [1]:

- *The digital image is only used as visual reference:* the required digital image quality is low, both in terms of spatial and color resolution content. The display is limited to a screen or a low-resolution print device. Exact color reproduction is not critical; the image has only to “resemble” the original artwork or scene. Images are usually compressed to save storage space, but primarily delivery time.
- *The digital image is used for print reproduction:* the requirements for the digitizing system depend on the desired reproduction quality. Limiting output to certain spatial dimensions and a certain color encoding will define the digitizing devices and initial processing. Commercial digital image archives and many cultural institutions digitize their analog image collections with this purpose in mind. Similarly, professional photographers who use digital cameras will build a collection of such images, and therefore an image database, over time
- *The digital image represents a “replacement” of the original in terms of spatial and color information content:* this goal is the most challenging to achieve. Color information can only be captured adequately if the digitizing device supports the originals’ color gamut in terms of encoding bit-depth and lack of metamerism. Such image library applications can be found in museums and other collections, where colorimetrically accurate capture is important, and/or the physical artwork is too fragile to be handled daily, or is already deteriorating so that the image content as seen today needs to be preserved.

Each of the usages outlined above will determine the quality of the *digital master* to be archived. It represents the highest quality file that has been digitized [2]. Since it represents the information that is supposed to survive long term, the encoding has to be appropriate for all current and future–as yet unknown–usage. The choices made in the initial digitization and processing of images are final and can usually not be reversed. For daily use, *derivatives* of digital master files can be created, which are encoded according to their purpose. In many image databases, various image derivatives are found, each appropriate for a different application [3]. See Table 1 for a comparison of different image database usage and recommended encoding specifications.

Table 1: Comparison of different image database purposes and recommended encoding specifications.

Master file	**Visual Representation**	**Print Reproduction**	**Replacement of the Original**
Purpose of the image database	• Consumer (digital camera files, print scanning) • On-line visual representation (electronic catalogues, image kiosks, insurance databases) • Low-end print representation (prints up to 4 x 6 inch)	• Commercial image libraries (picture agencies, publishing houses, news agencies) • Professional photographers • Museums, cultural, and government institutions • Consumer printing (8 x 10 inch)	• Museums, cultural, and government institutions with high quality demand and/or fragile original artwork.
Derivatives	• Screen resolutions • Thumbnails	• Lower print resolutions • Screen resolutions • Thumbnails	• Print resolutions • Screen resolutions • Thumbnails
Spatial Resolution[1] (for one dimension)	• Thumbnails: $\leq$ 250 pixels • Screen resolution: $\leq$ 1600 pixels	• Dependent on the reproduction intent, the size and quality of the original artwork, and/or limited by the maximum spatial resolution of the digitizing device. • ~ 2000 to 7000 pixels	• Dependent on the size and quality of the original artwork, and/or limited by the maximum spatial resolution of the digitizing device.
Image encoding of master file	Output-referred	Output-referred	Sensor Input-referred
Compression of the master file	None, lossless, or lossy (JPEG, JPEG2000)	None or lossless	None or lossless
File Format of the master file	EXIF, TIFF, JFIF, JP2	EXIF, TIFF, JPX	TIFF, JPX

[1]Note that pixel dimension is only one parameter describing resolution of a digital image file. Other factors, such as the size, resolution and sharpness of the original, the spatial frequency response of the digitizing system, the quality of pixel reconstruction and image processing algorithms, and the resolution of the output device all influence the resolution and apparent sharpness of the image reproduction.

2. DIGITAL IMAGE COLOR WORKFLOW

The color workflow of a digital image has been discussed before [4,5] and will only be briefly summarized here in relation to image databases. An image is captured and encoded into a sensor or source device space, which is device and image specific. It may then be transformed into an input-referred image representation, i.e. a color encoding describing the scene's or original's colorimetry. In most workflows, however, the image is directly encoded from source device space into an output-referred image representation, which describes the image appearance on some real or virtual output. If the output-referred color space describes a virtual output, then additional transforms are necessary to encode the image into output coordinates, which are dependent on the specific device (see Figure 1).

2.1 Sensor Encoding

When a scene or original is captured, either by a scanner or by a digital camera, its first color encoding is device and scene specific, defined by illumination, sensor, exposure, and filters. Images encoded in raw-data, or sensor space, are not viewable without further processing. However, they represent the true archive files necessary for many high-end image database applications. Any further processing can degrade the image, as algorithms are based on today's know-how and might be improved in the future. Additionally, these transforms are not always reversible, even when stored with the image files. In most cases, the hardware and software necessary to create and archive such master files is not commercially available and has to be developed in-house or with academic or industrial partners.

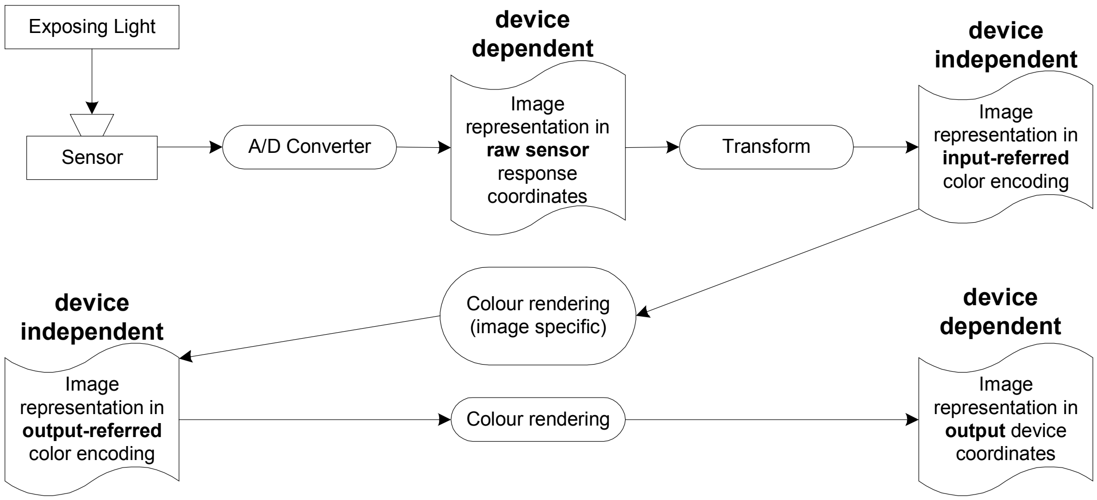

Figure 1: Color image workflow from capture to display.

2.2 Input-Referred Image Encoding

The purpose of input-referred image encoding is for the image to represent an estimate of the scene's or the original's colorimetry. An input-referred encoding maintains the relative dynamic range and gamut of the scene or original. A scene-referred encoding describes the input-referred encoding of an original scene, where as an original-referred encoding describes the input of original artwork. Scene-referred encodings usually need to support a higher dynamic range. The quality of the colorimetric estimate depends on the ability to choose the correct scene illuminant, and the correct transformations from device RGB to standard input-referred color encodings.

Color spaces that are used to encode an estimate of the scene's or original's colorimetry are usually CIE XYZ, CIE Lab, CIE and Luv. There are currently three input-referred color encodings that are well defined: Photo YCC [6], RIMM RGB [7,8], and the colorimetric ICC PCS [9]. Input-referred image encodings can be used for archiving master image files when it is important that the original colorimetry is preserved so that a facsimile can be created at a later date, and the transformation from sensor encoding to input-referred encoding is unambiguous. Another advantage of input-referred encodings, 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, without having to maintain all the device characterization data necessary to process a raw sensor data encoded image.

2.3 Output-Referred Image Encoding

Output-referred image encodings refer to image representations in color spaces that are based on the colorimetry of real or virtual output characteristics. Images can be transformed into output-referred encodings from either source or input-referred image encodings. The complexity of these transforms varies: they can range from a simple gamma-function, such as is employed by video encoding to complex image dependent algorithms. Output-referred color encodings, based on real or virtual output devices, are e-sRGB (e-sYCC) [10], sRGB [11], perceptual ICC PCS [9], ROMM RGB [8,12], Adobe 98 RGB [5], and SWOP CMYK [12]. Output-referred image representations can also be encoded in other color spaces, such as CIE Lab, CIE Luv, Photo YCC [6], YUV, and YCrCb.

Archiving images in output-referred encoding is appropriate for image libraries that archive primarily master files whose purpose is to be a reproduction of the original, and not a representation. Most current image libraries fall into this category, as commercial hardware and software are available for this task. The choice of which output-referred color encoding is appropriate for the master file and its derivatives depends on the database applications. However, whenever possible, the same encoding should be chosen for the master files, especially when large number of images are processed to derivatives at a later date. In general, all image library derivatives are encoded in output-referred representations.

2.4 Output Device Specific Encoding

Transforms from output-referred encodings to output coordinates are device and media specific. If an output-referred color space is equal or close enough to the real device characteristics, such as "monitor" RGBs, no additional transformation to device specific digital values is needed. For other output applications, such as print, there is a need for additional conversions.

Apart from some graphic arts applications, it is rare today that images are archived in output coordinates, such as device and media specific RGB, CMY, or CMYK spaces, as re-purposing of the image files for other applications and output devices is more complex. However, there are many legacy master 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.

3. INPUT- AND OUTPUT-REFERRED ENCODINGS

The correct color space and the correct color encoding parameters depend on the particular image application. The following parameters need to be considered [13]: extent of color gamut, perceptually linear tone-scale encoding to minimize the bit-depth needed to encode an image, dynamic range, illuminant white-point, viewing conditions, quantization efficiency, and compression. Table 2 summarizes the most important attributes of input referred and output-referred color encodings.

Table 2: Attributes of input-referred and output-referred color encodings.

	Input-referred encoding	**Output-referred encoding**
Image representation	Colorimetric estimate of a scene/original	Colorimetric estimate of a reproduction
Color gamut	Large enough to encompass most scene/original colors	Large enough to encompass most output devices
Perceptual uniformity (transfer function)	Data is optionally encoded using a transfer function for approximate perceptual uniformity (invertability desired)	Data is optionally encoded using a gamma-type power function to approximate perceptual uniformity on the output device (invertability desired)
Dynamic range	Must handle a scene luminance ratio of at least 10'000:1, or the luminance ratio of the original.	Must handle an image luminance ratio of at least 1000:1
White point	Should accommodate floating white points or chromatic adaptation to a fixed white point	Fixed white point determined by reproduction viewing conditions (D50, D65)
Viewing conditions (linkage to color appearance)	Luminance level, viewing surround, adapted white point, and viewing flare, as typical of outdoor environments for scene-referred encodings, typical of original viewing conditions for original encoding.	Luminance level, viewing surround, adapted white point, and viewing flare, as typical of indoor environments
Quantization/ Encoding	Quantization errors not visible on smooth, noiseless ramps Extended bit-depths encoding desired (10, 12, 16-bit per channel)	Quantization errors not visible on smooth, noiseless ramps, 8, 10, 12 or 16-bit encoding (8-bit for applications)
Compressibility	Not very important	Importance dependent on the imaging application (easy conversion to YCC color encoding)
Usual color encodings	CIEXYZ, CIELAB, CIELUV, RIMM RGB, Photo YCC	e-sRGB, sRGB, ROMM RGB, Adobe RGB 98, YCbCr, Photo YCC (legacy: Apple RGB)
Applications	Master files for high-end applications (replacement of the original)	Master files (for print reproduction): e-sRGB, ROMM RGB, Adobe RGB 98, Photo YCC Master files (for screen viewing): sRGB, YCbCr Derivatives for print reproduction or screen viewing

4. STANDARD COLOR ENCODINGS

4.1 Input-Referred Encodings

Many input-referred encodings are based on CIEXYZ, CIE Lab, or CIE Luv. That is especially appropriate where the input encoding is actually an original referred encoding, and the transformation from sensor to input-referred encoding is based on minimizing colorimetric errors. The encoding specifications, such as bit-depth, white-point, luminance ratio, viewing conditions, etc. depend on the application. In case of the colorimetric ICC PCS, the absolute intent defines a white-point luminance of 160 cd/m^2, where as for media relative, the media used determines the dynamic range. The white-point is D50, and the encoding bit-depth is variable (the ICC PCS defines either linear XYZ or non-linear Lab encoding).

Reference Input Medium Metric RGB (RIMM RGB) was proposed by Eastman Kodak and has been standardized by I3A [7]. The reference primaries describe a wide gamut RGB color space that allows possible color encodings outside of the spectral locus. The reference viewing conditions used to encode scene color values are typical of outdoors viewing environments. The observer adaptive white point has the chromaticity values of CIE standard illuminant D50, necessitating that a chromatic adaptation transformation must be applied to the image data in cases where the chromaticity of the observer adaptive white for an actual scene differs from that of the reference conditions. The non-linear transfer function is based on the ITU-R BT.709 recommendations and is equal to 0.45 (1/2.2). The luminance dynamic range can be encoded up to 200 percent of the exposure value associated with a normally exposed perfect diffuse white reflector in the scene. If that range is not sufficient, an extended range RGB (ERIMM RGB) is also defined that varies form RIMM RGB only in the non-linear encoding function. RIMM RGB master file encoding can be considered if a scene-referred input encoding is desired, or the dynamic range of the original is very high, and the corresponding output-referred derivative encoding is ROMM RGB. RIMM RGB allows for 8-, 12- and 16-bit encoding, ERIMM for 12- and 16-bit

Eastman Kodak has also defined Photo YCC [6], which is an encoding that was primarily developed for the *Kodak PhotoCD* scanning system. The specifications have been published, and many image libraries that have used – or still use – the PhotoCD system have master files archived in Photo YCC. Photo YCC was originally designed as an input-referred encoding for photographic slide and negative scanning, but can also be used as an output-referred color encoding. The viewing conditions are the same as RIMM RGB, but the scene adaptive white is assumed to be CIE illuminant D65. The encoding is based on a reference image capturing device that has the sensitivity of the ITU-R BT.709 color matching functions. These RGB709 values are non-linearly encoded with a gamma function of 0.45 (1 / 2.2) to R'G'B'709. R'G'B'709 is linearly transformed to Luma and chroma values, and the encoding is usually quantized to 8-bit. Photo YCC can easily be transformed to a monitor based RGB space for viewing, or to YC_bC_r for compression.

4.2 Output-referred Encodings

There are several output-referred color encodings that have been standardized, or have become de-facto industry standards. All have been designed with certain applications in mind. *sRGB* is defined by IEC 61966-2-1 [11] as a default color encoding for multimedia applications. The purpose of sRGB is to define an output-referred color encoding for data interchange in multimedia where ICC color management is not implemented. Due to similarities of the defined reference display to real CRT monitors, often no additional color space conversion is needed to display images. However, conversions are required to transform image data to devices with different dynamic ranges, gamuts and viewing conditions. sRGB has also become the de-facto encoding of consumer digital camera images. Most consumer desktop color printers today assume for the transformation to device specific coordinates that an image is encoded in sRGB unless otherwise specified. sRGB encoding is therefore ideal for all image library derivatives that involve either monitor viewing, such as publishing to the web or electronic kiosk applications, or consumer applications.

An extended gamut and bit-depth encoding, *e-sRGB* [10], has recently become an I3A standard. Most encoding parameters, such as transfer function, viewing conditions, and primaries are similar to sRGB. However, e-sRGB defines an offset and over range to encode a larger gamut and dynamic range than sRGB. The gamut of e-sRGB is very close to the gamut of all visible colors. e-sRGB is designed as a storage and interchange space for photographic applications, and is well suited to encode output-referred master files where photographic rendering is applied before archiving, but the output is unknown. Because of its similar encoding specifications, the conversion from e-sRGB to sRGB is simple to implement, so that derivative files intended for monitor viewing are easily created. Transformations to ROMM RGB, YCC and ICC PCS are also specified.

ROMM (Reference Output Medium Metric) RGB is a wide-gamut, output-referred RGB color encoding. 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 input-referred color space, such as RIMM RGB. ROMM RGB uses a CIE illuminant D50 white point, which is a standard for viewing and evaluating graphic arts reproductions, as well as the ICC PCS white point. 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. ROMM RGB defines therefore 12- and 16-bit encoding to allow for greater precision in addition to the usual 8-bit encoding. ROMM RGB encoded images do not display well without additional transformations to device specific values. The encoding parameters have rather been chosen to fit well into an ICC color managed graphic arts workflow. Therefore, ROMM RGB encoding is most appropriate for master files or derivative files that will be undergo further color transforms and/or printed on hard copy output.

Adobe, Inc. introduced the concept of a working space that is device independent in their currently most popular commercial image manipulation application, Adobe Photoshop©. The goal is to make the image data more portable and not tied to a RGB display. It is also the space the user will import images to from different sources and make editing decisions in. *Adobe RGB 98* [5] was intended to provide a larger gamut than previous monitor spaces so pre-press users can set it as the default working space in Adobe Photoshop. It is based on the SMPTE-240M standard and was later renamed Adobe RGB 98. The Adobe RGB 98 encoding gamut is smaller than the ROMM RGB gamut, but still encloses most device colors of current CMYK printers. It is close enough in gamut to halftone printing device gamuts that applying ICC color management usually results in good reproduction quality. It is therefore very popular in graphic arts and professional photography applications. Images encoded in Adobe RGB 98 also display reasonably well on an uncalibrated monitor, so that a reasonable good representation of image appearance can be achieved without color management. While the encoding is reasonable for derivative files that are distributed today to graphic arts applications and that might or might not be color managed, it is worth to consider archiving output-referred master files with a larger gamut so that more scene and original colors can be retained. It is foreseeable that future display and print applications will be able to reproduce a larger gamut and dynamic range.

Apple RGB [5] color encoding is based on the classic Apple 13" RGB monitor. Because of its popularity and similar Trinitron-based monitors that followed, many key publishing applications used it as the default RGB space in the past, and therefore many legacy images are encoded in this color space. The gamut of Apple RGB and sRGB are similar, but the non-linear transfer function differs (1/1.8 for Apple RGB, 1/ 2.4 for sRGB). Therefore, Apple RGB images will appear too dark on non-calibrated CRT monitors. This effect is especially visible in web applications where thumbnails with different color encodings are displayed together. However, it is not appropriate to transform Apple RGB master files to another output-referred color encoding, as the occurrence of visual quantization artifacts increases with the number of color transformations applied to an image, and it is likely that a new default encoding for display viewing will be defined once the pre-dominant display technologies change. Rather, new derivatives should be created for the application desired.[1]

4. SUMMARY

It is not possible today to give a general color encoding "recipe" for image databases, as the most suitable color encoding is strongly dependent on the purpose of the database, and these vary greatly. In theory, every image database can design its own color encodings optimized for the applications of master and derivative image files. The encoding would either be based on CIE XYZ, CIE Lab, CIE Luv, RGB, YCC or n-channel space. The encoding specifications need to be stored so that future processing – to new derivative files or to create new ICC profiles – can always be guaranteed. Care should be taken that the initial encoding can support the color transformation and processing steps that are necessary to create subsequent derivative files from the point of view of gamut, bit-depth and noise. However, while an application-specific encoding might be desirable for the master archive file, it is not usually reasonable to use non-standard image encodings for derivatives files. In practice, commercial applications drive most image production workflows and display, and are therefore the recipients of image database derivative files. They only understand a limited number of color encodings, and it is advantageous considering productivity and image quality to keep derivative files in a standard output color encoding. If a (higher quality) master file is available, new derivatives with different encodings for different output-driven purposes can always be generated from these archive files.

REFERENCES

1. F. Frey and S. Süsstrunk, "Image Quality Issues for the Digitization of Photographic Collections," *Proc. IS&T 49th Annual Conference,* pp. 349-353, 1996.
2. J. Holm & N. Judge: "Electronic Photography at the NASA Langley Research Center," *Proc. IS&T 48th Annual Conference*, pp. 436-441, 1995.

[1] Note that when transforming from Apple RGB to sRGB for monitor viewing, it is usually sufficient to just adjust the non-linear transfer function. The primaries are close enough so that in practical viewing applications, the slight color shift from one encoding to the next is not noticeable.

3. F.S. Frey and J.M. Reilly, *Digital Imaging for Photographic Collections*, Image Permanence Institute, RIT, 1999. Available at http://www.rit.edu/ipi.
4. J. Holm, "Integrating New Color Image Processing Techniques with Color Management*," Proc. IS&T/SID 7th Color Imaging Conference,* pp. 80-86, 1999.
5. S. Süsstrunk, R. Buckley, and S. Swen, "Standard RGB Color Spaces," *Proc. IS&T/SID 7th Color Imaging Conference,* pp. 127-134, 1999.
6. E.J. Giorgianni and T.E. Madden, *Digital Color Management – Encoding Solutions*, Addison-Wesley, 1998.
7. I3A 7466:2001 Photography – Electronic still picture imaging – *Reference Input Medium Metric RGB Color encoding (RIMM-RGB)*
8. K.E. Spaulding, G.J. Woolfe, E.J. Giorgianni, "Optimized Extended Gamut Color Encoding for Scene-Referred and Output-Referred Image States," *J. Imaging Science and Technology*, Vol. 45, pp. 418-426, 2001.
9. ICC.1:2001, *File Format for Color Profiles*, International Color Consortium.
10. I3A 7667: 2001 Photography – Electronic Still Picture Imaging – Extended sRGB color encoding: e-sRGB.
11. IEC 61966-2-1: 1999 Multimedia systems and equipment – *Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB.*
11. I3A 7666: 2001 Photography – *Electronic Still Picture Imaging – Reference Output Medium Metric RGB Color encoding: ROMM-RGB*
13. ISO/WD 22028-1.15 (2001): Photography and graphic technology – *Extended colour encoding for digital image storage, manipulation and interchange – Part 1: Architecture and Requirements.*