INTERNATIONAL STANDARD

ISO 22028-1

First edition
2004-03-15

Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange —

Part 1: Architecture and requirements

Photographie et technologie graphique — Codages par couleurs étendues pour stockage, manipulation et échange d'image numérique —

Partie 1: Architecture et exigences

Reference number
ISO 22028-1:2004(E)

PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

Contents

Page

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

ISO 22028-1 was prepared by Technical Committee ISO/TC 42, *Photography*, in collaboration with ISO/TC 130, *Graphic technology*, and the International Commission on Illumination (CIE).

ISO 22028 consists of the following parts, under the general title *Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange*:

— *Part 1: Architecture and requirements*

The following parts are under preparation:

— *Part 2: Reference output medium metric RGB colour image encoding (ROMM RGB)*

— *Part 3: Reference input medium metric RGB colour image encoding (RIMM RGB)*

Introduction

Modern digital imaging systems serve a variety of consumer and commercial applications. Depending on the application, differing priorities will apply to such system attributes as image quality, interoperability, simplicity of system architecture and computations, and the flexibility for optimally using images for a variety of purposes. Trade-offs among these attributes are application-dependent.

A fundamental choice for any imaging system architecture is how to represent images numerically, in what colour space and with what digital encoding. In some applications, a single colour encoding designed to be compatible with the prevalent mode of image viewing by the end-user may suffice. Since both multimedia and internet-based imaging rely heavily on the viewing of images on a softcopy display, the use of sRGB as a colour encoding makes sense for those applications. However, because the colour gamut of sRGB does not encompass the colour gamuts of many common input and output devices, a system architecture that depends exclusively on the use of sRGB would compromise colour reproduction accuracy unacceptably for some applications.

Colour management systems, such as that defined by the International Color Consortium (ICC), provide a mechanism for transforming between various device-dependent and device-independent colour encodings through the use of colour profiles that are used to define transformations between the various colour encodings and a standard colour space known as the profile connection space (PCS). (The ICC.1:2001-12 specification defines two different PCS variations; one for colorimetric intent profiles, and one for perceptual intent profiles.) The ICC PCS is intended to be a colour space to be used for connecting together different colour profiles, and as such has a colour gamut large enough to encompass most common input and output devices and media. However, the ICC PCS was not designed to be used as a colour encoding for the storage, transmission or editing of digital images. Additionally, since ICC colour management is primarily designed to work with colour images in a picture-referred image state, it does not provide any explicit mechanism for the representation and manipulation of image data corresponding to other image states.

There are many different applications in the fields of digital photography and graphic technology that involve editing, storage and interchange of digital images in a variety of image states and colour encodings. In order to clearly communicate colour image information within and between these applications, it is necessary to unambiguously describe the meaning of the colour values used to encode digital images. The colour encoding definitions need to not only include a specification of the relationship between the digital code values and corresponding physical colour values, but they also need to clearly specify any other information needed to unambiguously interpret the colour values. Accordingly, there is a need to identify what information is required when defining a colour encoding in order to ensure that digital image data can be clearly communicated between various applications.

This part of ISO 22028 addresses this need by specifying a set of requirements to be met by colour encodings defined for various digital imaging applications. This part of ISO 22028 also describes a reference image-state-based digital imaging architecture that is flexible enough to support a wide variety of applications and workflows. This image-state-based digital imaging architecture can be used to classify colour encodings into a number of different image states. However, this part of ISO 22028 does not specify any particular workflow(s) that need to be used for any particular digital imaging applications.

There is also a need for the specification of standard extended-gamut colour encodings that can be used in the context of this architecture to preserve the full range of colour information at every stage of the workflow, from the initial image capture through to the final step of producing a softcopy or hardcopy reproduction. It is anticipated that subsequent parts of this multi-part standard will define at least one scene-referred extended-gamut colour encoding and at least one output-referred extended-gamut colour encoding.

The International Organization for Standardization (ISO) draws attention to the fact that it is claimed that compliance with this document may involve the use of a patent concerning colour management given in Clauses 4 and 5.4.3.

ISO takes no position concerning the evidence, validity and scope of this patent right.

The holder of this patent right has assured ISO that he/she is willing to negotiate licences under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of this patent right is registered with ISO. Information may be obtained from:

Microsoft Corporation

One Microsoft Way, Redmond, WA 98052.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights other than those identified above. ISO shall not be held responsible for identifying any or all such patent rights.

Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange —

Part 1: Architecture and requirements

1 Scope

This part of ISO 22028 specifies a set of requirements to be met by any extended-gamut colour encoding that is to be used for digital photography and/or graphic technology applications involving digital image storage, manipulation and/or interchange. This part of ISO 22028 is applicable to pictorial digital images that originate from an original scene, as well as digital images with content such as text, line art, vector graphics and other forms of original artwork. This part of ISO 22028 also describes a reference image-state-based digital imaging architecture, encompassing many common workflows, that can be used to classify extended colour encodings into a number of different image states. However, this part of ISO 22028 does not specify any particular workflow(s) that are to be used for digital photography and/or graphic technology applications.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/CIE 10527, *CIE standard colorimetric observers*

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

3.1
absolute colorimetric coordinates
tristimulus values, or other colorimetric coordinates derived from a tristimulus values, where the numerical values correspond to the magnitude of the physical stimulus

EXAMPLE When CIE 1931 standard colour-matching functions are used, the Y-coordinate value corresponds to the luminance, not the luminance factor (or some scaled value thereof).

3.2
adapted white
colour stimulus that an observer who is adapted to the viewing environment would judge to be perfectly achromatic and to have a luminance factor of unity; i.e. absolute colorimetric coordinates that an observer would consider to be a perfect white diffuser

NOTE The adapted white may vary within a scene.

3.3
additive RGB colour space
colorimetric colour space having three colour primaries (generally red, green and blue) such that CIE XYZ tristimulus values can be determined from the RGB colour space values by forming a weighted combination of

the CIE XYZ tristimulus values for the individual colour primaries, where the weights are proportional to the radiometrically linear colour space values for the corresponding colour primaries

NOTE 1 A simple linear 3 × 3 matrix transformation can be used to transform between CIE XYZ tristimulus values and the radiometrically linear colour space values for an additive RGB colour space.

NOTE 2 Additive RGB colour spaces are defined by specifying the CIE chromaticity values for a set of additive RGB primaries and a colour space white point, together with a colour component transfer function.

3.4
adopted white
spectral radiance distribution as seen by an image capture or measurement device and converted to colour signals that are considered to be perfectly achromatic and to have an observer adaptive luminance factor of unity; i.e. colour signals that are considered to correspond to a perfect white diffuser

NOTE 1 The adopted white may vary within a scene.

NOTE 2 No assumptions should be made concerning the relation between the adapted or adopted white and measurements of near perfectly reflecting diffusers in a scene, because measurements of such diffusers will depend on the illumination and viewing geometry, and other elements in the scene that may affect perception. It is easy to arrange conditions for which a near perfectly reflecting diffuser will appear to be grey or coloured.

3.5
colorimetric colour space
colour space having an exact and simple relationship to CIE colorimetric values

NOTE Colorimetric colour spaces include those defined by CIE (e.g. CIE XYZ, CIELAB, CIELUV, etc.), as well as colour spaces that are simple transformations of those colour spaces (e.g. additive RGB colour spaces).

3.6
colour component transfer function
single variable, monotonic mathematical function applied individually to one or more colour channels of a colour space

NOTE 1 Colour component transfer functions are frequently used to account for the nonlinear response of a reference device and/or to improve the visual uniformity of a colour space.

NOTE 2 Generally, colour component transfer functions will be nonlinear functions such as a power-law (i.e. “gamma”) function or a logarithmic function. However, in some cases a linear colour component transfer function may be used.

3.7
colour encoding
generic term for a quantized digital encoding of a colour space, encompassing both colour space encodings and colour image encodings

3.8
colour gamut
solid in a colour space, consisting of all those colours that are either: present in a specific scene, artwork, photograph, photomechanical, or other reproduction; or capable of being created using a particular output device and/or medium

3.9
colour image encoding
digital encoding of the colour values for a digital image, including the specification of a colour space encoding, together with any information necessary to properly interpret the colour values such as the image state, the intended image viewing environment and the reference medium

NOTE 1 In some cases the intended image viewing environment will be explicitly defined for the colour image encoding. In other cases, the intended image viewing environment may be specified on an image-by-image basis using metadata associated with the digital image.

NOTE 2 Some colour image encodings will indicate particular reference medium characteristics, such as a reflection print with a specified density range. In other cases the reference medium will be not applicable, such as with a scene-referred colour image encoding, or will be specified using image metadata.

NOTE 3 Colour image encodings are not limited to pictorial digital images that originate from an original scene, but are also applicable to digital images with content such as text, line art, vector graphics and other forms of original artwork.

3.10
colour-matching functions
tristimulus values of monochromatic stimuli of equal radiant power

[CIE Publication 17.4, 845-03-23]

3.11
colour rendering
mapping of image data representing the colour-space coordinates of the elements of a scene to output-referred image data representing the colour-space coordinates of the elements of a reproduction

NOTE Colour rendering generally consists of one or more of the following: compensating for differences in the input and output viewing conditions, tone scale and gamut mapping to map the scene colours onto the dynamic range and colour gamut of the reproduction, and applying preference adjustments.

3.12
colour re-rendering
mapping of picture-referred image data appropriate for one specified real or virtual imaging medium and viewing conditions to picture-referred image data appropriate for a different real or virtual imaging medium and/or viewing conditions

NOTE Colour re-rendering generally consists of one or more of the following: compensating for differences in the viewing conditions, compensating for differences in the dynamic range and/or colour gamut of the imaging media, and applying preference adjustments.

3.13
colour space
geometric representation of colours in space, usually of three dimensions

[CIE Publication 17.4, 845-03-25]

3.14
colour space encoding
digital encoding of a colour space, including the specification of a digital encoding method, and a colour space value range

NOTE Multiple colour space encodings may be defined based on a single colour space where the different colour space encodings have different digital encoding methods and/or colour space value ranges. (For example, 8-bit sRGB and 10-bit e-sRGB are different colour space encodings based on a particular additive RGB colour space.)

3.15
colour space white point
colour stimulus to which colour space values are normalized

NOTE The colour space white point may or may not correspond to the assumed adapted white point and/or the reference medium white point for a colour image encoding.

3.16
continuous colour space value
real-valued, unbounded colour space value that has not been encoded using a digital encoding method

3.17
device-dependent colour space
colour space defined by the characteristics of a real or idealized imaging device

NOTE Device-dependent colour spaces having a simple functional relationship to CIE colorimetry can also be categorized as colorimetric colour spaces. For example, additive RGB colour spaces corresponding to real or idealized CRT displays can be treated as colorimetric colour spaces.

3.18
digital imaging system
system that records and/or produces images using digital data

3.19
extended gamut
colour gamut extending outside that of the standard sRGB CRT display as defined by IEC 61966-2-1

3.20
film rendering transform
mapping of image data representing measurements of a photographic negative to output-referred image data representing the colour-space coordinates of the elements of a reproduction

3.21
film unrendering transform
mapping of image data representing measurements of a photographic negative to scene-referred image data representing estimates of the colour-space coordinates of the elements of the original scene

3.22
gamut mapping
mapping of the colour-space coordinates of the elements of a source image to colour-space coordinates of the elements of a reproduction to compensate for differences in the source and output medium colour gamut capability

NOTE The term “gamut mapping” is somewhat more restrictive than the term “colour-rendering” because gamut mapping is performed on colorimetry that has already been adjusted to compensate for viewing condition differences and viewer preferences, although these processing operations are frequently combined in reproduction and preferred reproduction models.

3.23
hardcopy
representation of an image on a substrate which is self-sustaining and reasonably permanent

[ISO 3664]

3.24
ICC profile
International Color Consortium's file format, used to store transforms from one colour encoding to another, e.g. from device colour coordinates to profile connection space, as part of a colour management system

3.25
image state
attribute of a colour image encoding indicating the rendering state of the image data

NOTE The primary image states defined in this document are the scene-referred image state, the original-referred image state and the output-referred image state.

3.26
International Color Consortium profile connection space (ICC PCS)
standard colour image encoding defined by the International Color Consortium providing a standard connection point for combining ICC profiles

NOTE The ICC.1:2001 specification defines two variations of the PCS, an original-referred variation for colorimetric intent profiles, and an output-referred variation for perceptual intent profiles.

3.27

luminance factor

ratio of the luminance of the surface element in the given direction to that of a perfect reflecting or transmitting diffuser identically illuminated.

[CIE Publication 17.4, 845-04-69]

3.28

luminance ratio

ratio of the maximum luminance to the minimum luminance that is either: present in a specific scene, artwork, photograph, photomechanical, or other reproduction; or is capable of being created using a particular output device and medium

3.29

medium black point

neutral colour with the lowest luminance that can be produced by an imaging medium in normal use, measured using the specified measurement geometry

NOTE It is generally desirable to specify a medium black point that has the same chromaticity as the medium white point.

3.30

medium white point

neutral colour with the highest luminance that can be produced by an imaging medium in normal use, measured using the specified measurement geometry

3.31

metadata

data associated with a digital image aside from the pixel values that comprise the digital image

NOTE Metadata are typically stored as tags in the digital image file.

3.32

original-referred image state

image state associated with image data that represents the colour-space coordinates of the elements of a two-dimensional hardcopy or softcopy image, typically produced by scanning artwork, photographic transparencies or prints, or photomechanical or other reproductions

NOTE 1 When the phrase “original-referred” is used as a qualifier to an object, it implies that the object is in an original-referred image state. For example, original-referred image data is image data in an original-referred image state.

NOTE 2 Original-referred image data are related to the colour-space coordinates of the original, typically measured according to ISO 13655, and do not include any additional veiling glare or other flare.

NOTE 3 The characteristics of original-referred image data that most generally distinguish them from scene-referred image data are that they refer to a two-dimensional surface, and the illumination incident on the two-dimensional surface is assumed to be uniform (or the image data corrected for any non-uniformity in the illumination).

NOTE 4 There are classes of originals that produce original-referred image data with different characteristics. Examples include various types of artwork, photographic prints, photographic transparencies, emissive displays, etc. When selecting a colour re-rendering algorithm, it is usually necessary to know the class of the original in order to determine the appropriate colour re-rendering to be applied. For example, a colorimetric intent is generally applied to artwork, while different perceptual algorithms are applied to produce photographic prints from transparencies, or newsprint reproductions from photographic prints. In some cases the assumed viewing conditions are also different between the original classes, such as between photographic prints and transparencies, and will usually be considered in well-designed systems.

NOTE 5 In a few cases, it may be desirable to introduce slight colorimetric errors in the production of original-referred image data, for example to make the gamut of the original more closely fit the colour space, or because of the way the image data were captured (such as a Status A densitometry-based scanner).

3.33
output-referred image state
image state associated with image data that represents the colour-space coordinates of the elements of an image that has undergone colour-rendering appropriate for a specified real or virtual output device and viewing conditions

NOTE 1 When the phrase “output-referred” is used as a qualifier to an object, it implies that the object is in an output-referred image state. For example, output-referred image data are image data in an output-referred image state.

NOTE 2 Output-referred image data are referred to the specified output device and viewing conditions. A single scene can be colour-rendered to a variety of output-referred representations depending on the anticipated output-viewing conditions, media limitations, and/or artistic intents.

NOTE 3 Output-referred image data may become the starting point for a subsequent reproduction process. For example, sRGB output-referred image data are frequently considered to be the starting point for the colour re-rendering performed by a printer designed to receive sRGB image data.

3.34
picture-referred image state
image state associated with image data that represents the colour-space coordinates of the elements of a hardcopy or softcopy image, encompassing both original-referred image data and output-referred image data

NOTE 1 When the phrase “picture-referred” is used as a qualifier to an object, it implies that the object is in a picture-referred image state. For example, picture-referred image data are image data in a picture-referred image state.

NOTE 2 Picture-referred image data will generally be colour-rendered for a specific real or virtual imaging medium and viewing condition.

NOTE 3 Picture-referred image data can include image data that do not originate from an original scene, such as text, line art, vector graphics and other forms of original artwork.

3.35
scene
spectral radiances of a view of the natural world as measured from a specified vantage point in space and at a specified time

NOTE A scene may correspond to an actual view of the natural world or to a computer-generated virtual scene simulating such a view.

3.36
scene-referred image state
image state associated with image data that represents estimates of the colour-space coordinates of the elements of a scene

NOTE 1 When the phrase “scene-referred” is used as a qualifier to an object, it implies that the object is in a scene-referred image state. For example, scene-referred image data are image data in a scene-referred image state.

NOTE 2 Scene-referred image data can be determined from raw DSC image data before colour-rendering is performed. Generally, DSCs do not write scene-referred image data in image files, but some may do so in a special mode intended for this purpose. Typically, DSCs write standard output-referred image data where colour-rendering has already been performed.

NOTE 3 Scene-referred image data typically represent relative scene colorimetry estimates. Absolute scene colorimetry estimates may be calculated using a scaling factor. The scaling factor can be derived from additional information such as the image OECF, FNumber or ApertureValue, and ExposureTime or ShutterSpeedValue tags.

NOTE 4 Scene-referred image data may contain inaccuracies due to the dynamic range limitations of the capture device, noise from various sources, quantization, optical blurring and flare that are not corrected for, and colour analysis errors due to capture device metamerism. In some cases, these sources of inaccuracy can be significant.

NOTE 5 The transformation from raw DSC image data to scene-referred image data depends on the relative adopted whites selected for the scene and the colour space used to encode the image data. If the chosen scene adopted white is inappropriate, additional errors will be introduced into the scene-referred image data. These errors may be correctable if the transformation used to produce the scene-referred image data is known, and the colour encoding used for the incorrect scene-referred image data has adequate precision and dynamic range.

NOTE 6 The scene may correspond to an actual view of the natural world, or may be a computer-generated virtual scene simulating such a view. It may also correspond to a modified scene determined by applying modifications to an original scene to produce some different desired scene. Any such scene modifications should leave the image in a scene-referred image state, and should be done in the context of an expected colour-rendering transform.

3.37
softcopy
representation of an image produced using a device capable of directly representing different digital images in succession and in a non-permanent form

EXAMPLE The most common example is a monitor.

[ISO 3664]

3.38
standard original-referred colour encoding
a colour encoding for original-referred image data defined and documented by an authorized standards body or industry consortium

3.39
standard output-referred colour image encoding
a colour image encoding for output-referred image data defined and documented by an authorized standards body or industry consortium

3.40
standard scene-referred colour image encoding
a colour image encoding for scene-referred image data defined and documented by an authorized standards body or industry consortium

3.41
tristimulus value
amounts of the three reference colour stimuli, in a given trichromatic system, required to match the colour of the stimulus considered

[CIE Publication 17.4, 845-03-22]

3.42
veiling glare
light, reflected from an imaging medium, that has not been modulated by the means used to produce the image

NOTE 1 Veiling glare lightens and reduces the contrast of the darker parts of an image.

NOTE 2 In CIE 122, the veiling glare of a CRT display is referred to as ambient flare.

3.43
viewing flare
veiling glare that is observed in a viewing environment but not accounted for in radiometric measurements made using a prescribed measurement geometry

NOTE The viewing flare is expressed as a percentage of the luminance of adapted white.

4 Image-state-based digital imaging architecture

4.1 General

The architecture of a digital imaging system can be described, on the one hand, as the sum of its components and how those components are interconnected and, on the other hand, as the functions of those components and how they interact with each other as an integrated system. One important aspect of a digital imaging architecture is how the digital image data is encoded as it progresses through the system workflow from image capture/creation, through image processing/storage/interchange, and finally to output on one or more output devices.

The need for various colour 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. The digital photography and graphic technology industries are very diverse and often complex. However, their core activities can be represented by a fairly simple model where images are classified according to their image state. As shown in Figure 1, this model consists of a generic digital imaging architecture that can be used to describe the workflows for many different applications. Examples showing how a number of typical workflows can be described in the context of this architecture are given in Annex A.

This image-state-based digital imaging architecture, and the associated terminology, facilitates a common framework for classifying different colour encodings, and describing imaging chains for many diverse types of digital imaging systems. Any colour image encodings that are defined in this multi-part standard shall be described within the context of this architecture. In particular, the colour image encoding shall be identified with an appropriate image state. This part of ISO 22028 does not specify any workflows that should be used for any particular applications to transform image data to/from the identified image states.

The digital imaging architecture shows examples of where different types of devices may fit within typical workflows utilizing colour image encodings compliant with this part of ISO 22028. It is not intended to constrain the workflows for any particular applications to those shown in Figure 1. For example, raw digital camera captures may be processed directly to an output-referred colour encoding without stopping in a scene-referred colour encoding.

Workflows associated with particular applications may include additional colour encodings that may correspond to image states different than the standard image states defined in this image-state-based digital imaging architecture. For example, it may be useful to define a colour encoding for representing colour negative scans, or an intermediate colour encoding for partially colour-rendered images. While such colour encodings may be valuable internally to particular applications, they should generally not be used for image interchange in an open system environment unless all components in the system are enabled to properly interpret/use the image data and/or provision is made for communicating a transformation of the image data to one of the standard image states.

The image state diagram shown in Figure 1 shows that most colour encodings can broadly be categorized into scene-referred or picture-referred image states.

4.2 Scene-referred colour encodings

Scene-referred colour encodings are representations of the estimated colour-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.

EXAMPLE Scene-referred image data may be represented in many different ways including encoding scene colour values using a CIE colour space such as CIE XYZ or CIELAB, or in terms of the response of an idealized scene capture device such as RIMM RGB.

Scene-referred image data may correspond to an actual view of the natural world, or to a computer-generated virtual scene simulating 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 over-exposed regions of a back-lit 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 colour-rendering transform. For example, typical colour-rendering transforms will include a boost in the chroma of the image. Any boost in colourfulness of the scene (e.g. making the grass greener) should be done with the knowledge that there will be an additional chroma boost during colour-rendering. Consequently, the

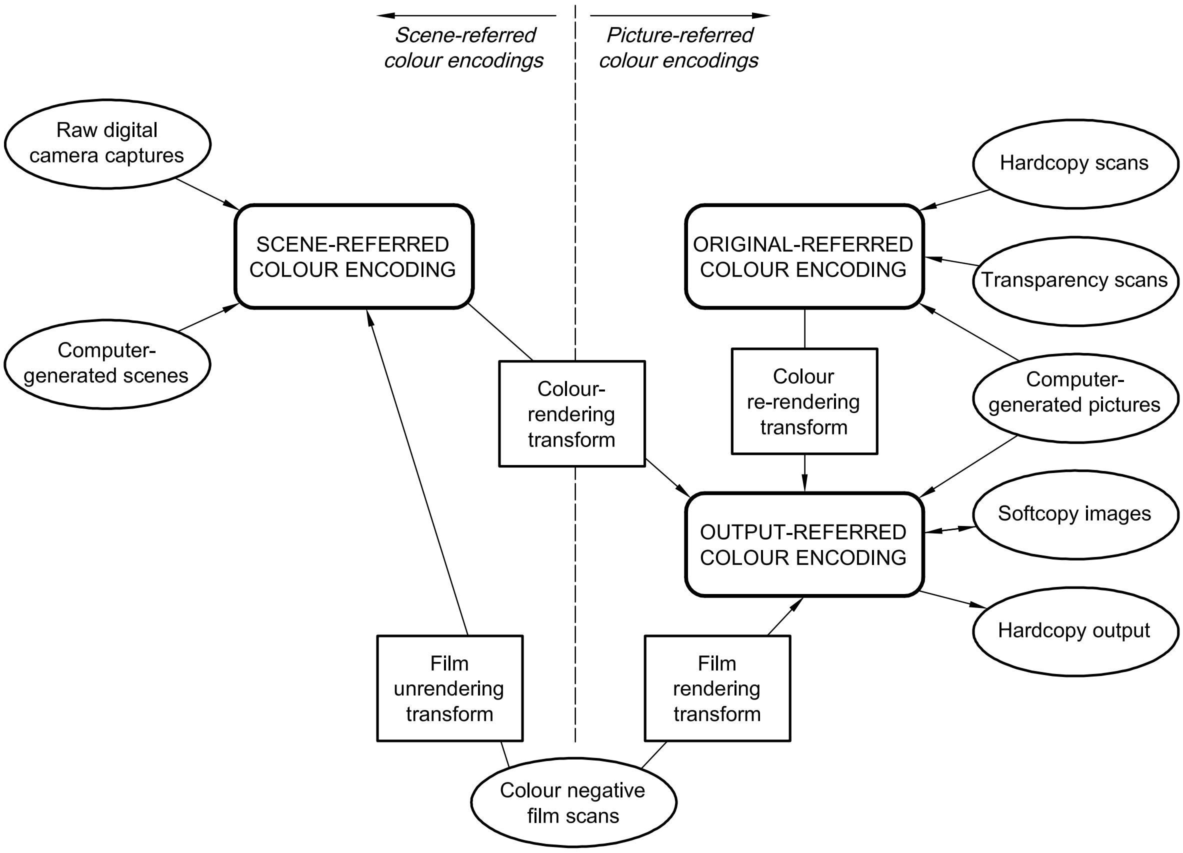

Figure 1 — Image state diagram showing relationship between various types of colour encodings

colour-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. Image modifications that change the image state of the image (e.g. tone scale and gamut-mapping operations that colour-render an image to map the scene colours onto the dynamic range and colour gamut of a certain output medium) are inappropriate for determining a modified scene.

It should be noted that the image colorimetry of the scene-referred image data may contain inaccuracies due to the dynamic range limitations of the capture device, noise from various sources, quantization, optical blurring and flare that are not corrected for, and colour analysis errors due to capture device metamerism. In some cases, these sources of inaccuracy can be significant.

NOTE Since scene-referred image data have not been colour-rendered, they are usually not ready to be either displayed or printed.

4.3 Picture-referred colour encodings

4.3.1 General

Picture-referred colour encodings are representations of the colour-space coordinates of a hardcopy or softcopy image. Picture-referred colour encodings can be further subdivided into original-referred colour encodings and output-referred colour encodings.

4.3.2 Original-referred colour encodings

Original-referred colour encodings are representative of the colour-space coordinates (or an approximation thereof) of a two-dimensional hardcopy or softcopy input image. Original-referred image data are 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.

NOTE Examples of ways that original-referred image data may be represented include encoding picture colour values using a CIE colour space such as CIE XYZ or CIELAB, 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 scanning device.

When a scene-capture device, such as a digital camera, is used for the purposes of digitizing a two-dimensional hardcopy or softcopy image, the resulting image data should generally be treated as original-referred image data rather than scene-referred image data. In this case, it is usually unnecessary and inappropriate to apply a colour-rendering transform to the image data for purposes of determining output-referred image data since the original image has already been colour-rendered. However, it may be desirable to apply a colour re-rendering transform to account for the differences between the media/viewing condition characteristics of the original image source and the final output-referred image.

4.3.3 Output-referred colour encodings

Output-referred colour encodings are representative of the colour-space coordinates of image data that are appropriate for a specified real or virtual output device and viewing conditions. Output-referred colour encodings are tightly coupled to the characteristics of a particular real or virtual output device and viewing conditions. Standard output-referred colour encodings are most commonly used for image data intended for open interchange.

NOTE 1 Examples of ways that output-referred image data may be represented include encoding picture colour values using a CIE colour space such as CIE XYZ or CIELAB, by a colour encoding derived from CIE colorimetry (e.g. sRGB or ROMM RGB), or by device-dependent control signals for a particular softcopy or hardcopy output device.

NOTE 2 In some cases, output-referred image data can become the starting point for a subsequent reproduction process. For example, sRGB standard output-referred image data are frequently considered to be the starting point for many desktop printing systems. In this case, the printing systems will generally perform a colour re-rendering process to transform the sRGB colour values to those appropriate for the particular output device and assumed viewing conditions.

4.4 Colour-rendering transforms

A colour-rendering transform is used to transform a scene-referred image to an output-referred image. Colour-rendering transforms embody the tone and colour reproduction aims of an imaging system, relating the corresponding scene colorimetry to the desired picture colorimetry.

It should be noted that colour-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 colour gamut of the output medium, and applying colour adjustments to account for preferences of human observers.

NOTE Colour-rendering transforms are typically proprietary and irreversible.

4.5 Colour re-rendering transforms

For cases where original-referred image data are to be transformed to output-referred image data, a colour re-rendering transform should generally be used to adjust the image colorimetry when the media and/or intended viewing condition characteristics are not the same.

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 colour encoding associated with a reference reflection print medium and viewing conditions, then the colour 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 colour gamut.

NOTE 1 Colour re-rendering transforms are typically proprietary and can be irreversible depending on the relative dynamic range and colour gamut of the original-referred and output-referred image data.

Since most real input/output media and viewing conditions will not be exact matches of those associated with a particular output-referred (or original-referred) colour encoding, some colour re-rendering is usually necessary in order to map the actual device colour values to/from a particular output-referred (or original-referred) colour encoding. In some cases, the colour re-rendering transform may be as simple as gamut mapping any colour values outside the gamut of a particular output device. In other cases, more complex colour re-rendering transforms may be used that include modification of colours inside the device colour gamut and/or compensation for viewing condition differences, etc.

NOTE 2 In imaging systems employing ICC colour management, colour re-rendering to account for specific input/output device characteristics is typically incorporated into the perceptual and saturation intents in ICC profiles for the devices. ICC colour management systems make provision for supporting several re-rendering transform options by defining a number of standard rendering intents. The ICC colorimetric rendering intents are designed to accurately preserve in-gamut colours. The ICC perceptual rendering intent is vendor-specific and involves compromises such as trading off preservation of contrast in order to preserve detail throughout the tonal range. It is intended to be used for the general reproduction of images, particularly pictorial or photographic-type images. The ICC saturation rendering intent is vendor-specific and involves compromises such as trading off preservation of hue in order to preserve the vividness of pure colours. It is intended to be useful for images that contain objects such as charts or diagrams. Since the ICC PCS version 4 perceptual-intent reference medium colorimetry represents an output-referred colour encoding, perceptual intent transforms in input profiles for devices producing scene-referred image data (e.g. raw digital camera captures) should also incorporate any colour-rendering needed to produce an appropriate rendered picture.

4.6 Film rendering and unrendering transforms

Colour negative film scans can be transformed to either a scene-referred image state or an output-referred image state. Due to 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 colours from the film scan. Alternatively, the film scan can be used to directly form output-referred image data. 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.

In some cases, colour negative film may be used to photograph a two-dimensional hardcopy or softcopy input image. In this case, applying a film unrendering transform to the colour negative film scans would produce original-referred image data rather than scene-referred image data.

5 Requirements for specifying a colour encoding

5.1 Colour encoding hierarchy

One critical component of a digital imaging system architecture is the specification of how digital images are represented numerically, and what metadata are associated with the digital images to further define the interpretation of the image data. First, it is necessary to define the colour space in which the image will be encoded. It is also necessary to define a quantized digital colour encoding of the colour space. Colour encodings can be specified at two levels. A colour space encoding includes the specification of a digital encoding method. A colour image encoding further includes any additional information necessary to properly interpret the image colour values such as the image state, the image viewing environment and the reference imaging medium. Figure 2 illustrates this hierarchical relationship between a colour space, a colour space encoding and a colour image encoding.

NOTE 1 According to this terminology, many “colour spaces” that have been defined in various standards and other documents are more correctly referred to as “colour space encodings” or “colour imaging encodings.” For example, an ICC version 4 perceptual intent “colour space” (which includes the specification of the image state, image viewing environment and reference imaging medium) would be a “colour image encoding” according to the terminology used in this part of ISO 22028. Similarly, the sRGB “colour space” would also be an example of a “colour image encoding.”

Colour encodings defined for the purpose of communicating colour image information in open digital imaging systems should generally be complete colour image encodings in order to ensure unambiguous communication of the intended image characteristics. Colour space encodings or incomplete colour image encodings should only be used for communicating colour image data in cases where provision is made for conveying any

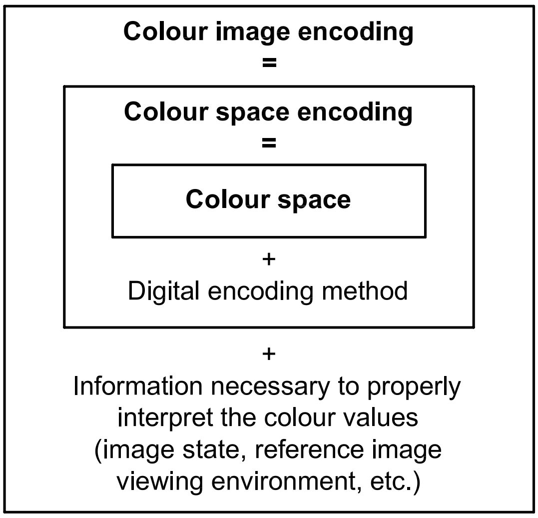

Figure 2 — Hierarchical relationship between colour space, colour space encoding and colour image encoding

remaining information needed to define a complete colour image encoding using mechanisms such as image file metadata, or in cases where the colour image data are only used within a closed digital imaging system where such information is well understood throughout the system.

In cases where image file metadata are used to convey any remaining information needed to define a complete colour image encoding, a file format supporting the storage of appropriate metadata should be used. This solution is only appropriate where it is known that any applications intended to utilize the digital image will be enabled to properly understand and use the metadata.

Even in cases where a completely defined colour image encoding is used to represent a digital image, provision should be made for identifying the particular colour image encoding to all system components. When digital images are exchanged in an open system, such identification can be made through the use of image file metadata provided that all system components are enabled to properly understand and use the metadata.

NOTE 2 Some file formats require the storage of images in a particular standard colour image encoding, or specify that the image data can be assumed to be in such a standard colour image encoding by default. In such cases, it is possible to infer the colour image encoding without explicitly identifying it using metadata.

Subclause 5.2 further describes the information needed to define a colour space; 5.3 further describes the information needed to define a colour space encoding; and 5.4 further describes the information needed to define a colour image encoding. Annex B summarizes the characteristics of a number of existing standard colour image encodings in accordance with the terminology specified in this part of ISO 22028. Annex C discusses a number of criteria that should be considered when selecting a colour image encoding for use in a particular application and workflow.

5.2 Information needed to define a colour space

5.2.1 General

The information needed to define a colour space will depend on the type of colour space being specified. Colour spaces can be broadly categorized into three types:

— colorimetric,

— colour appearance, and

— device-dependent.

In some cases, a nonlinearity, or some other colour transformation, may be applied to the colour space values for purposes such as improving encoding efficiency. The resulting transformed colour space values shall be considered to be in a different colour space.

EXAMPLE If a logarithmic transformation were applied to CIE XYZ colour values, the original XYZ colour values and the log XYZ colour values would be two different colour spaces.

5.2.2 Colorimetric colour spaces

5.2.2.1 General

For colorimetric colour spaces, the relationship between the continuous colour space values and CIE colorimetry shall be defined. In some cases, the colour space may be one of the standard CIE colour spaces such as CIE XYZ or CIELAB. In those cases, the colour spaces are themselves measures of CIE colorimetry. In other cases, a colour space may be specified that has a defined relationship to CIE colorimetry. In those cases, the transformation back and forth between the colour space and standard CIE colorimetry shall be defined. A CIE standard colorimetric observer, as defined in ISO/CIE 10527, shall be used for the computation of all colorimetric quantities. The particular standard colorimetric observer that is used shall be specified.

NOTE The CIE 1931 standard colorimetric observer (often referred to as the 2° standard observer) will be used in most applications.

Since colorimetric quantities will be dependent on measurement conditions, the assumed measurement conditions associated with all CIE colorimetry shall be specified (e.g. 0/45 geometry versus integrating sphere). For cases where the colorimetric quantities are computed from reflective or transmissive media, the illuminant shall also be specified.

5.2.2.2 CIE colour spaces

5.2.2.2.1 CIE colour space specification

A number of colorimetric colour spaces have been defined by the CIE including CIE XYZ, CIELAB and CIELUV. These colour spaces can be used directly for encoding digital image data. In such cases, the appropriate CIE colour space specification shall be identified.

5.2.2.2.2 Colour space white point

The definition of a CIE colour space shall include the specification of the CIE chromaticity values and the luminance value for the colour space white point.

5.2.2.3 Additive RGB colour spaces

5.2.2.3.1 General

Additive RGB colour spaces are one common type of colorimetric colour space. They are defined by specifying a set of additive RGB primaries, a colour space white point, and a colour component transfer function.

5.2.2.3.2 RGB primaries chromaticity values

The definition of an additive RGB colour space shall include the specification of the CIE chromaticity values for the RGB primaries.

NOTE The RGB primaries generally correspond to colours that are in the red, green and blue regions of colour space. However, this is not a requirement, and in fact some sets of primaries could contain physically unrealizable colours outside of the spectrum locus.

5.2.2.3.3 Colour space white point

The definition of an additive RGB colour space shall include the specification of the CIE chromaticity values and the luminance value for the colour space white point. A phosphor matrix determined from the CIE chromaticity values of the RGB primaries and the colour space white point can be used to relate the linear RGB colour space values to the corresponding CIE tristimulus values.

NOTE When the colour-matching functions associated with the RGB primaries are suitable as cone fundamentals for von Kries-type chromatic adaptation, the RGB colour space white point can be conveniently mapped to multiple white points in CIE XYZ. Such an additive RGB colour space can be used to achieve a level of white point independence, and multiple valid transformations to different XYZ colour space white points can be provided.

5.2.2.3.4 Colour component transfer function

The definition of an additive RGB colour space shall include the specification of a forward colour component transfer function that shall be applied to the radiometrically linear RGB colour space values to determine transformed RGB colour space values. Radiometrically-linear colour space values are linear with respect to image radiance. An inverse colour component transfer function shall also be specified for transforming the RGB colour space values back to corresponding radiometrically-linear RGB colour space values. Either the forward or inverse colour component transfer function shall be identified to be normative, the other one being derived by inverting the normative transform. Preferably, the inverse colour component transfer function should be normative.

NOTE Generally, the colour component transfer function will be a nonlinear function such as a power-law (i.e. “gamma”) function or a logarithmic function. However, in some cases a linear colour component transfer function could be used.

5.2.2.4 Luma-chroma colour spaces derived from additive RGB colour spaces

5.2.2.4.1 General

Luma-chroma colour spaces are another common class of colorimetric colour spaces. Luma-chroma colour spaces are derived from additive RGB colour spaces using a luma-chroma matrix transformation applied to the RGB colour space values.

NOTE Luma-chroma colour spaces are sometimes generically referred to as YCC or YC_RC_B or YC_BC_R colour spaces.

5.2.2.4.2 Additive RGB colour space

The definition of a luma-chroma colour space shall include the specification of an additive RGB colour space upon which the luma-chroma colour space is based. (See 5.2.2.3.)

5.2.2.4.3 Luma-chroma matrix

The definition of a luma-chroma colour space shall include the specification of a forward luma-chroma matrix that shall be applied to the RGB colour space values to determine luma-chroma colour space values. An inverse luma-chroma matrix shall also be specified for transforming the luma-chroma colour space values back to corresponding RGB colour space values. Either the forward or inverse luma-chroma matrix shall be identified to be normative, the other one being derived by inverting the normative matrix. Preferably, the inverse luma-chroma matrix should be normative.

NOTE The resulting “luminance-chrominance” values generally are not true “luminance” or “chrominance” values, but are only approximate correlates to these colorimetric quantities.

5.2.2.5 Other colorimetric colour spaces

Other forms of colorimetric colour spaces can also be defined that do not fall into one of the classes enumerated here. The relationship between the continuous colour space values and the corresponding image colorimetry for any such colour space shall be clearly defined.

5.2.3 Colour appearance colour spaces

Colour spaces based on colour appearance models, such as CIECAM97s, can be used for encoding images in certain applications. The definition of a colour appearance colour space shall include the specification of the particular colour appearance model upon which the colour space is based, together with the particular colour coordinates to be used to represent the image data.

NOTE Colour appearance colour spaces are generally derived from CIE colorimetry, together with additional parameters relating to the image viewing environment.

5.2.4 Device-dependent colour spaces

5.2.4.1 General

Some colour 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 colour spaces. Input-device-dependent colour spaces are defined by the response of a reference image capture device. Output-device-dependent colour spaces are defined by the characteristics of a reference output device.

NOTE 1 Examples of device-dependent colour spaces would include densitometric input-device colour spaces such as standard Status A densities, and output-device colour spaces such as printer CMY(K).

NOTE 2 Device-dependent colour spaces commonly occur within colour processing workflows but are rarely used for interchange of digital image data in open systems unless the characteristics of the reference device are clearly understood across all system components or provision is made to associate device characterization data with the digital image data (e.g. using an ICC profile).

5.2.4.2 Input-device-dependent colour spaces

5.2.4.2.1 General

Input-device-dependent colour spaces shall be defined by specifying the characteristics of a reference image capture device, including the spectral sensitivity and colour component transfer function. The reference image capture device may correspond to an actual image capture device, or to some idealized image capture device. A colour space white point shall also be specified.

NOTE 1 Input-device-dependent colour spaces can sometimes be related to CIE colorimetry if the spectral characteristics of the input media are known. For example, if a densitometer is used to measure photographic material where the spectral characteristics of the dye set and base material are known, it is possible to infer the image spectra, and from there determine the corresponding CIE colorimetry. However, the exact relationship between an input-device-dependent colour space and CIE colorimetry generally will be media-dependent, and therefore impossible to determine precisely for unknown input media.

NOTE 2 In many cases, the reference image capture device will have three colour channels — typically red, green and blue — however, this is not a requirement. For example, a spectral image capture device could be defined for capturing image data on a wavelength-by-wavelength basis.

NOTE 3 In some cases, it is inappropriate to use a simple model based on a specified spectral sensitivity and colour component transfer function to represent the response of some input devices. (For example, some real input devices might include internal colour transformations to colour spaces other than the fundamental sensor response space.) In such cases, it would be more appropriate to define the colour space by specifying the relationship between the continuous colour space values and the corresponding image colorimetry as described in 5.2.2.5.

5.2.4.2.2 Spectral sensitivity

The definition of an input-device-dependent colour space shall include the specification of the spectral sensitivity of the reference image capture device.

Input-device-dependent colour spaces based on image capture devices where the CIE colour-matching functions are expressible as linear combinations of the spectral sensitivities will have a direct relationship to CIE colorimetry and should be categorized as colorimetric colour spaces.

The spectral sensitivity for input-device-dependent colour spaces based on image capture devices having defined illumination sources should be defined by the spectral product of the sensor spectral sensitivities and the illumination source spectral power. For example, the effective spectral sensitivity for a film scanner would be the spectral product of the light source used to illuminate the film and the actual sensor spectral sensitivities. When fluorescence effects contribute significantly to an image capture process, the spectral composition of the illumination source should be specified in addition to the spectral product.

5.2.4.2.3 Colour component transfer function

The definition of an input-device-dependent colour space shall include the specification of a forward colour component transfer function for each colour channel of the colour space that relates the radiometrically-linear representation of the colour space values to the corresponding device-dependent colour space values. (Radiometrically-linear colour space values are linear with respect to image radiance.) An inverse colour component transfer function shall also be specified for transforming the device-dependent colour space values back to corresponding radiometrically-linear representation of the colour space values. Either the forward or inverse colour component transfer function shall be identified to be normative, the other one being derived by inverting the normative transform. Preferably, the inverse colour component transfer function should be normative.

NOTE Generally, the colour component transfer function will be a nonlinear function such as a power-law (i.e. “gamma”) function or a logarithmic function. However, in some cases a linear colour component transfer function could be used.

5.2.4.2.4 Colour space white point

The definition of an input-device-dependent colour space shall include the specification of a colour space white point. When the input-device-dependent colour space is intended to be used for representing images captured under a variety of illuminants, the colour space white point will generally correspond to an adopted white point.

5.2.4.3 Output-device-dependent colour spaces

Output-device-dependent colour space shall be defined by specifying the relationship between the control signals of a reference output device and the corresponding output image colour. The output image colour can be expressed in a number of different ways including output image spectra, output image colorimetry or output image density values. The reference output device may correspond to an actual output device, or to some idealized output device. The definition of such a space shall include the specification of a colour space white point.

Output-device-dependent colour spaces defined by a simple functional relationship to CIE colorimetry should be categorized as colorimetric colour spaces. For example, additive RGB colour spaces corresponding to real or idealized CRT displays can be treated as colorimetric colour spaces as described in 5.2.2.3.

Output-device-dependent colour spaces having incomplete or inexact relationships to CIE colorimetry should generally not be categorized as colorimetric colour spaces.

EXAMPLE A colour space defined by a sampled colour characterization data set enumerating the CIE colorimetry for a limited set of colour values will typically not be categorized as a colorimetric colour space.

NOTE In many cases the reference output device will have three colour channels — typically red, green and blue, or cyan, magenta and yellow — however, this is not a requirement. For example, many hardcopy output devices include an additional black colour channel.

5.3 Information needed to define a colour space encoding

5.3.1 General

Defining a colour space encoding requires the identification of a colour space as well as a digital encoding method.

5.3.2 Colour space

The definition of a colour space encoding shall include the specification of a colour space according to the guidelines described in 5.2.

5.3.3 Digital encoding method

5.3.3.1 General

The definition of a colour space encoding shall include the definition of a digital encoding method for each colour component specifying the relationship between the continuous colour 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.

Both a forward encoding transform (indicating how the continuous colour space values are mapped to the digital representation of those values) and an inverse encoding transform (indicating how the digital colour space values are mapped back to the continuous colour space values) shall be specified. Either the forward or inverse encoding transform shall be identified to be normative, the other one being derived by inverting the normative encoding transform. Preferably, the inverse encoding transform should be normative.

5.3.3.2 Integer digital encoding

5.3.3.2.1 General

The definition of an integer digital encoding method shall include the specification of a colour space value range and a corresponding digital code value range. The digital code values shall be linearly related to the continuous colour space values. If a nonlinear encoding of the colour space values is desired, the appropriate nonlinearity shall be incorporated into the definition of the colour space so that a linear quantization function can be used.

5.3.3.2.2 Colour space value range

The definition of an integer digital encoding method shall include the specification of a colour space value range defined by minimum and maximum continuous colour space values to be represented in the integer digital encoding.

NOTE 1 The continuous colour space values and the corresponding colour space value range will generally be in the natural units of the colour space, and will typically not be scaled to correspond to any particular integer range. For example, an RGB colour space might have a colour space value range of 0,0 to 1,0, and a CIELAB colour space might have colour space value ranges of 0,0 to 100,0 and –150,0 to 150,0 for the lightness and chroma colour channels, respectively.

NOTE 2 The minimum and maximum continuous colour space values are independent of any black and white points specified for the colour space. In some cases they might coincide, but in others they might not.

5.3.3.2.3 Digital code value range

The definition of an integer digital encoding method shall include the specification of a digital code value range defined by minimum and maximum integer digital code values, I_{min} and I_{max}, corresponding to the minimum and maximum continuous colour space values. (The number of digital code values in an integer digital encoding may be computed as $I_{\text{max}} - I_{\text{min}} + 1$.)

Typically the minimum continuous colour space value will be mapped to the minimum integer digital code value, and the maximum continuous colour space value will be mapped to the maximum integer digital code value. However, in some cases it may be desirable to invert the polarity so that the minimum continuous colour space value will be mapped to the maximum integer digital code value and the maximum continuous colour space value will be mapped to the minimum integer digital code value. In such cases the colour space encoding specification shall explicitly note this fact.

NOTE Often, the number of digital code values will be given by 2^N, where N is the bit-depth of the integer digital encoding.

For some applications, it is useful to define a family of colour space encodings having different digital code value ranges, but based on a single colour space and colour space value range. This enables the use of bit-depth independent colour transforms in colour-managed digital imaging systems. Such families of colour space encodings can be identified by a single name. When a single name is used to refer to a family of colour space encodings, the colour space encodings shall differ only by a scale factor.

A family of integer colour space encodings based on a single colour space and having a particular colour space value range are preferably specified using quantization equations that support a range of digital code value ranges; however, only a particular set of digital code value ranges may be supported in a given application.

The digital code value range is independent of the bit-depth of the digital memory used to store a digital image. For example, a digital image represented using a “12-bit” colour space encoding having a digital code value range of 0 to 4 095 may be stored in a 16-bit integer memory buffer. In cases where there is a mismatch between the bit-depth of the colour space encoding and the bit-depth of the digital memory, it is important that any application using the digital image data be aware of the true digital code value range of the colour space encoding and how it is stored in the digital memory.

5.3.3.3 Floating-point digital encoding

The definition of a floating-point digital encoding method shall include the specification of the relationship between the floating-point colour space values and the continuous colour space values.

5.3.4 Set of valid colour values

There may be certain subsets of colour values in the colour space encoding that are defined to be invalid and therefore should not be used for representing images stored in the colour space encoding. If any such invalid colour values exist, the colour space encoding specification shall identify the set of valid colour values.

The set of valid colour values can be specified relative to the encoded colour space values or relative to the continuous colour space values. The specification shall clearly identify which type of colour values are used to define the set of valid colour values.

NOTE Examples of colour values that might be specified to be invalid would include colour values that correspond to physically unrealizable colours (e.g. colours having negative CIE XYZ tristimulus values or having CIE chromaticity values outside the spectrum locus).

5.4 Information needed to define a colour image encoding

5.4.1 General

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

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

For some applications, it is useful to define a family of colour image encodings based on a corresponding family of colour space encodings that differ only in the digital code value range. This enables the use of bit-depth independent colour transforms in colour-managed digital imaging systems. Such families of colour image encodings can be identified according to a single name. When a single name is used to refer to a family of colour image encodings, the colour image encodings in the family shall differ only by a scale factor; all of the other colour image encoding attributes (e.g. image state or reference image viewing environment) shall be held constant.

For some applications, it is useful to define multiple colour image encodings that differ only in the colour space encoding, but which share a common image state and reference viewing environment, etc. Examples of this are sRGB and sYCC. and ROMM RGB and ICC v4 perceptual intent PCS. In such cases, simple colour space transformations can be used to convert between the colour image encodings. For example, a simple luma-chroma matrix can be used to convert between sRGB and sYCC. However, if two colour image encodings differ in image state and/or reference viewing environment, etc., any transformation between such colour image encodings should account for those differences.

EXAMPLE To transform from a scene-referred colour image encoding to an output-referred colour image encoding, a colour-rendering transformation needs to be applied to map the scene colour values to appropriate picture colour values. (This would be true even in the case where the scene-referred and output-referred colour image encodings are based upon the same colour space.)

Applications and document file formats that are designed to work with multiple colour image encodings shall be aware of any differences in image state, etc., and deal with them accordingly, or should limit themselves to working with colour image encodings that have these attributes in common.

In some cases, information such as the reference image viewing environment and reference imaging medium characteristics can be defined for a particular colour image encoding so that it is fixed for all images to be stored in that colour 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.

5.4.2 Colour space encoding

The definition of a colour image encoding shall include the specification of a colour space encoding according to the guidelines described in 5.3.

5.4.3 Image state

The definition of a colour image encoding shall include the specification of the image state to be associated with images stored in that colour image encoding. Where possible, the image state should be identified to be a scene-referred image state, an original-referred image state or an output-referred image state according to the descriptions given in Clause 4.

Certain colour image encodings (e.g. densitometric colour space encodings intended to be used for scanning colour negatives) may not fit within one of the standard image states described in Clause 4. If such image data are to be used in open exchange, all information needed to unambiguously interpret image data in the new image state shall be defined and thoroughly described, and all applications that are anticipated to use the image data should be enabled to properly interpret and use image data in the new image state.

The image state of certain colour image encodings may be somewhat ambiguous. For example, colour encodings used for colour video systems can be associated with the colorimetry of an original scene using the encoding equations for standardized video capture systems. They can also be associated with the colorimetry of an output picture using the characteristics of a standard display device. (This basically involves using two different colour image encodings, one relating the colour values to the scene colorimetry, and a second relating the colour values to the picture colorimetry.) In such cases, the image state identified for the colour image encoding shall be consistent with the image colorimetry specified in the colour image encoding definition. (For example, if the colour values are related to original scene colorimetry, the colour image encoding would be associated with a scene-referred image state; if the colour values are related to output picture colorimetry, the colour image encoding would be associated with an output-referred image state.) In cases where defined relationships may exist between the colour image encoding and image colorimetry in multiple image states, the colour image encoding should be identified with the image state that is most consistent with the fundamental characteristics of the colour encoding.

5.4.4 Reference image viewing environment

5.4.4.1 General

In order to properly interpret the colour appearance of images encoded in a colour image encoding, it is necessary to specify a reference image viewing environment in which the image is intended to be viewed.

NOTE It is not necessary that the image actually be viewed in the reference image viewing environment. Rather, the reference image viewing environment is intended to provide a context for interpreting the intended colour appearance of the encoded image colorimetry. For cases where the image is to be viewed in an actual viewing environment significantly different than the specified reference image viewing environment, it might be desirable to use a colour appearance transform to determine corresponding image colorimetry that would produce the intended colour appearance in the actual viewing environment. Depending on the application, it might not always be necessary or desirable to apply such colour appearance transforms, even for cases where the reference and actual viewing environments might be significantly different.

For colour appearance colour spaces, the definition of a reference image viewing environment is not required since it is accounted for in the determination of the colour values. However, for certain applications, it still may be useful to specify a reference image viewing environment in order to communicate that an image has been rendered for a certain viewing environment.

5.4.4.2 Image surround

The definition of a colour image encoding shall include the specification of an image surround. The image surround refers to the characteristics of the field surrounding the image being viewed, filling the field of vision. Where possible, the reference image surround should be quantitatively defined by specifying the surround chromaticity and surround luminance. The surround luminance can either be specified in absolute terms, or as some fraction of the colour space white point luminance or the adapted white point luminance.

In some applications the assumption of a uniform surround field is inappropriate. In these cases the distribution of light in the surround region should be carefully described.

5.4.4.3 Assumed adapted white point

The definition of a colour image encoding should include the specification of the CIE chromaticity values and the luminance value for an assumed adapted white point, the adapted white point being the colour stimulus that an observer who is adapted to the viewing environment would judge to be perfectly achromatic and to have a luminance factor of unity (i.e. absolute colorimetric coordinates that an observer would consider to be a perfect white diffuser).

If a colour image encoding is intended to be used for representing images where no assumed adapted white point can be identified, this fact should be clearly identified. Examples of applications where such encodings could be useful would include digital cinema applications where the adapted white point may be image-dependent and may vary on a scene-by-scene basis to achieve a desired visual effect, and applications where scene-referred images are encoded prior to any adapted white luminance estimation step.

NOTE 1 Depending on the image presentation, a particular medium white point might be perceived to be brighter than, darker than, or equal to the adapted white point. For example, by controlling the image presentation, the white point of a CRT might be perceived to be identical to a white piece of paper (i.e. somewhat darker than the adapted white point), or it might be perceived to have a luminance factor greater than unity (i.e. somewhat brighter than the adapted white point).

NOTE 2 The specification of the adapted white point for an output-referred colour image encoding effectively defines the expected appearance of the reference medium white point, and therefore has implications for the expected image presentation as well as the intended appearance of images to be stored using that colour image encoding.

NOTE 3 In a real image, the actual adapted white point might be different than the assumed adapted white point, and might even vary within the image.

NOTE 4 No assumptions should be made concerning the relation between the assumed adapted white point and measurements of near perfectly reflecting diffusers in a scene or picture, because measurements of such diffusers will depend on the illumination and viewing geometry, and other elements in the scene that might affect perception. It is easy to arrange conditions for which a near perfectly reflecting diffuser will appear to be grey or coloured.

5.4.4.4 Luminance of adapting field

The definition of a colour image encoding should include the specification of the luminance of the adapting field. This quantity can be specified by directly specifying a luminance value (typically in units of cd/m^2). Alternatively, the luminance of the adapting field can be defined relative to the adapted white point. For many typical viewing environments, setting the luminance of the adapting field to be 20 % of the luminance of the adapted white point is an appropriate default.

NOTE The luminance of the adapting field is used by many colour appearance models, such as CIECAM97s.

5.4.4.5 Viewing flare

The definition of a colour image encoding shall include the specification of the assumed viewing flare associated with the reference image viewing environment. The viewing flare should be specified as a percentage of the luminance of the colour space white point luminance or the assumed adapted white point luminance.

NOTE Viewing flare does not include any portion of the veiling glare that is accounted for in the measured image colorimetry (and therefore in the colour space values). If the measurement technique specified in the colour space definition exactly simulates the reference image viewing environment, then the viewing flare is identically zero.

5.4.5 Set of valid colour values

There may be certain subsets of colour values in the colour space encoding that are defined to be invalid and therefore should not be used for representing images stored in the colour image encoding. If any such invalid colour values exist, the colour image encoding specification shall identify the set of valid colour values.

The set of valid colour values can be specified relative to the encoded colour space values or relative to the continuous colour space values. The specification shall clearly identify which type of colour values are used to define the set of valid colour values.

NOTE Examples of colour values that might be specified to be invalid would include colour values that correspond to physically unrealisable colours (e.g. colours having negative CIE XYZ tristimulus values or having CIE chromaticity values outside the spectrum locus), or colour values outside of a specified dynamic range and/or reference colour gamut associated with the colour image encoding.

If a set of valid colour values is identified for the colour space encoding, the set of valid colour values for the colour image encoding should be a subset of those identified for the colour space encoding.

The set of valid colour values may be different for various colour image encodings that are based on the same colour space.

5.4.6 Reference imaging medium

5.4.6.1 General

The definition of a colour image encoding for an output-referred image state shall define the characteristics of a reference imaging medium associated with the specified real or virtual output device.

NOTE 1 In some cases, it might also be desirable to define a reference imaging medium for original-referred colour encodings.

Measurement conditions shall be specified for all colorimetric quantities associated with the definition of the reference imaging medium. The measurement conditions should be consistent with those specified for any colorimetric quantities associated with the colour space.

NOTE 2 All colorimetric quantities will therefore include any veiling glare associated with the specified measurement conditions.

5.4.6.2 Reference imaging medium white point

The definition of an output-referred colour image encoding shall include the specification of the CIE chromaticity values and luminance value for the reference imaging medium white point.

NOTE The reference imaging medium white point does not necessarily correspond to either the colour space white point or the assumed adapted white point.

5.4.6.3 Reference imaging medium black point

The definition of an output-referred colour image encoding shall include the specification of the CIE chromaticity values and luminance value for the reference imaging medium black point.

5.4.6.4 Rendering target colour gamut

The definition of an output-referred colour image encoding shall include the specification of any rendering target colour gamut associated with the reference imaging medium.

If a rendering target colour gamut is defined, it should be specified whether the colour gamut represents a firm boundary such that any colour values outside of that boundary are invalid, or whether the colour gamut is defined only to provide guidance for any colour-rendering operations that are used to prepare images to be stored in the colour image encoding.

A rendering target colour gamut may be defined corresponding to the actual colour gamut of a reference display device. Alternatively, a virtual colour gamut may be defined encompassing the colour gamut of one or more real or imaginary display devices.

Annex A
(informative)

Example system workflows

A.1 General

To illustrate how the image-state-based digital imaging architecture described in Clause 4 can be used in the context of many different applications, this annex gives examples of a number of different workflows commonly encountered in digital photography and graphic technology. These workflows take various paths through the image state diagram that was shown in Figure 1.

A.2 Generic workflow for digital photography

Figure A.1 illustrates a generic workflow for digital photography. The input image consists of the scene being photographed as seen through the optics of a digital still camera. Next, the image sensor and its associated electronics convert the optical image into a digital image representing the device-dependent raw response of the camera. A camera compensation transformation is then applied to the image in order to determine colour values in a scene-referred colour encoding. This transformation may include operations such as dark current and flare removal, exposure adjustment and white-balance, as well as a sensor characterization transform relating the raw sensor response to an estimate of the apparent scene colorimetry.

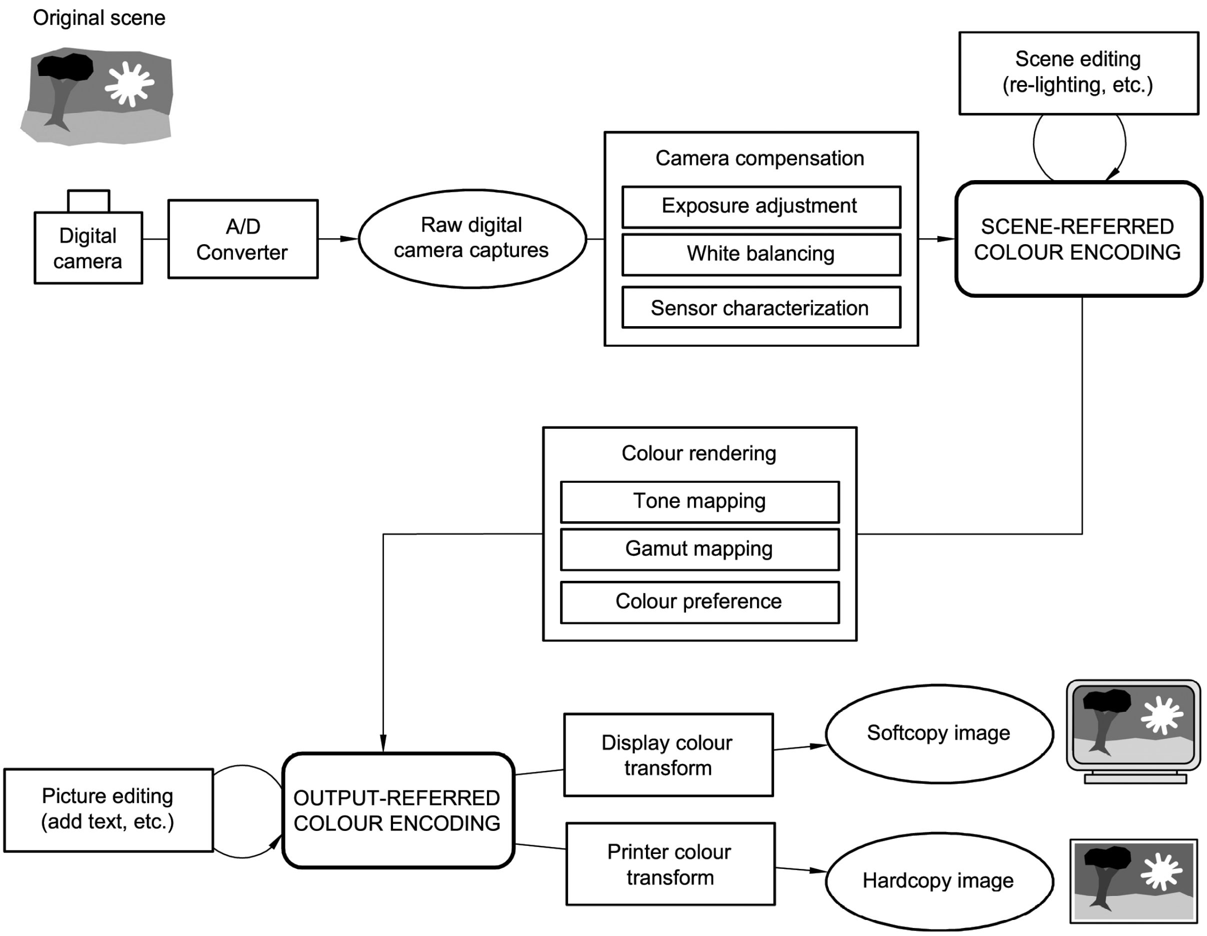

Figure A.1 — Generic workflow for digital photography

Generally, some sort of error minimization technique is used to determine the camera compensation transformation that results in minimum average colour errors under some set of nominal scene capture conditions. (Standard error minimization techniques for determining digital still camera compensation transformations are currently being developed as ISO 17321-2.) The accuracy of the camera compensation transformation will be a function of the sensor spectral sensitivities and the scene illumination, as well as the accuracy and robustness of the error minimization technique used to determine the camera compensation transformation. Consequently, the accuracy of the scene-referred image data will also depend on these same factors. The accuracy requirements for the scene-referred image data will be application-dependent.

At this point, the image data represent an estimate of the original scene colorimetry. However, in some cases, it may be desirable to modify the scene-referred image data to produce a preferred scene. This could include operations such as 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 over-exposed regions of a back-lit scene. (This can be viewed as being analogous to “re-lighting” the scene.) The scene-editing process 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 scene-editing operations should leave the image in a scene-referred image state, and should be done in the context of the expected colour-rendering transform. For example, typical colour-rendering transforms will include a boost in the chroma of the image. Any boost in colourfulness of the scene (e.g. making the grass greener) should be done with the knowledge that there will be an additional chroma boost during colour-rendering. Consequently, the colour-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.

Next, a colour-rendering transform is applied. This operation applies the tone/colour reproduction aims for the imaging system and changes the state of the image from a scene-referred image state to an output-referred image state. Colour-rendering transforms 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 colour gamut of the output medium, and applying colour preference adjustments to account for colour reproduction preferences of human observers. (Colour-rendering transforms are typically proprietary and irreversible.)

Additional image-editing operations may be applied to the output-referred image data. For example, the image may be modified using an interactive image-editing application, text or graphics may be added to the image, the image may be cropped/rotated/resized, or the image may be composited with other images.

Finally, the image is transformed to a device-dependent representation for a specific output device using a device colour transform, and a final viewable image is produced on the output device as either a softcopy or hardcopy image. The device colour transform relates the colour space values in the output-referred colour encoding to the appropriate device control signals necessary to reproduce the image on that device. The device colour transform may include a re-rendering transform to account for any differences between the dynamic ranges and colour gamuts of the reference output device and the actual output device.

The following should be noted.

— In some workflows, the device colour transform may be implemented using ICC profiles for the output-referred colour encoding and the output device.

— In some applications, the output-referred colour encoding may correspond to a colour encoding for a particular reference output device, such as an sRGB display. In that case, no device colour transform would be needed to send the image to the reference output device.

— An output-referred colour encoding produced by a digital photography workflow may be the starting point for a subsequent graphic technology workflow.

In some workflows, the various processing steps shown in Figure A.1 may be distributed, involving a different set of persons/applications for each step of the process that may or may not directly communicate with each other. The person who originally captured the image may not know for what purpose an image will be used, or at least not know what output device will be used. In this case, the greatest flexibility can be retained by deferring colour-rendering decisions as long as possible in the workflow.

Depending on the details of the particular workflow, the end users may or may not have any control over the various transformation steps, or even be aware that they are occurring. For example, in a professional workflow the camera compensation transform may be applied by the digital camera's in-camera processing and the resulting scene-referred image data stored in a proprietary file format that can only be opened using custom image acquisition software. The colour-rendering transform may be applied in the image acquisition software, which may or may not provide opportunity for the customer to adjust the colour-rendering aims. The image acquisition software may then provide the output-referred image data to a commercial image-editing software in an extended-gamut colour encoding for further editing. The image-editing software would apply an appropriate display colour transform to provide a preview image during the image-editing process. Finally, a colour-managed print path could be used to print the final image. In the other extreme, for a fully-automatic consumer workflow, all of the colour transforms could be applied in the digital camera's in-camera processing and the resulting image could be stored directly in a display-ready colour encoding with no opportunity for user interaction or control.

A.3 Generic workflow for digital camera producing display-ready images

Some of the intermediate image representations shown in Figure A.1 may not be explicitly employed in every digital photography workflow. Depending on the particular application, some (or all) of the image-processing transforms may be concatenated together. In the extreme, a single colour transform operation may be used to go directly from the raw digital camera capture through to the final output device colour encoding so that there are no intermediate images at all. For example, Figure A.2 shows a workflow that is common in many consumer digital photography applications. In this case, the digital camera produces a finished file in an output-referred, display-ready colour encoding (e.g. sRGB). The in-camera processing converts the raw image captures all the way through to the final display-ready colour encoding. There may or may not be intermediate stopping points in the image-processing chain, but if there are they are usually not accessible to the user. Since the images produced by the digital camera are in a display-ready colour encoding, they can be displayed directly on a softcopy display without the need for additional colour transforms (provided that the characteristics of the softcopy display are sufficiently close to those of the standard display assumed for the display-ready colour encoding). A printer colour transform is still required to print hardcopy images, but in most desktop digital imaging applications this colour transform is applied in the printer driver or hardware and is invisible to the user.

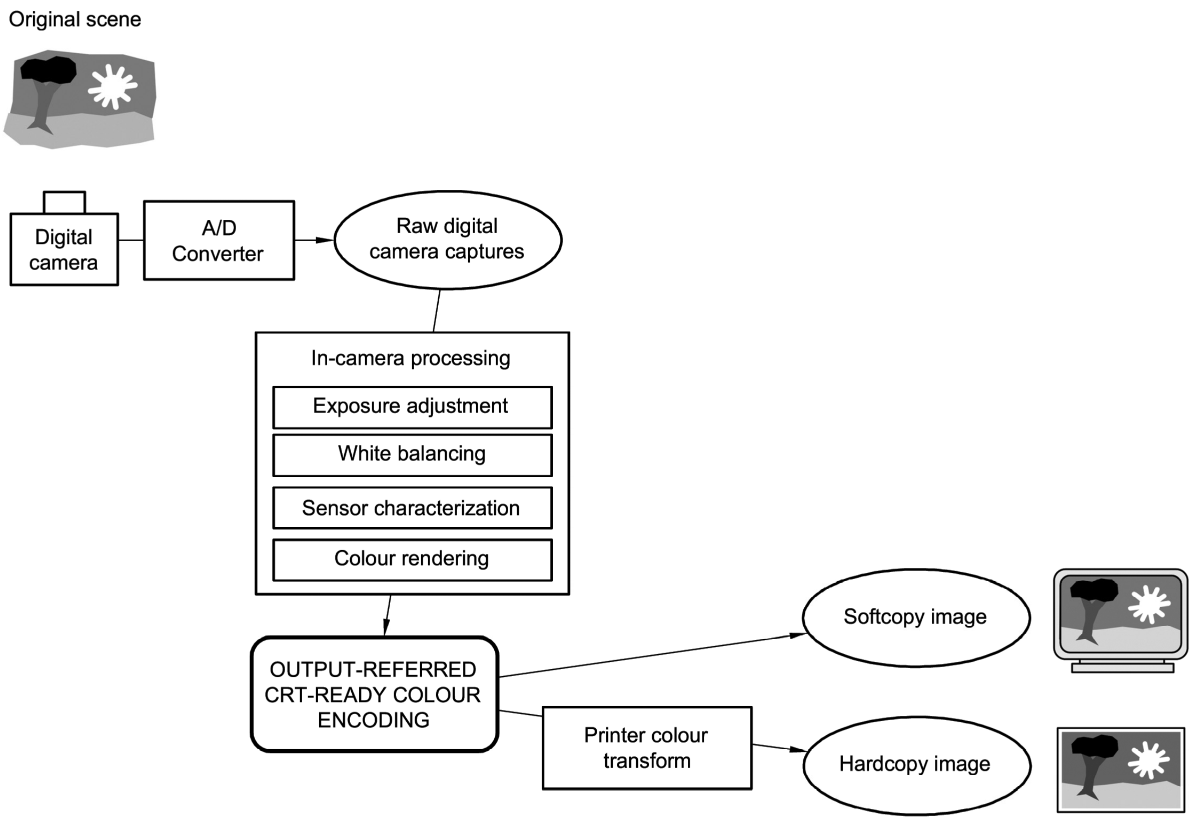

Figure A.2 — Generic workflow for digital camera producing display-ready images

The in-camera processing will typically include operations such as exposure adjustment and white-balance, as well as any camera compensation and colour-rendering necessary to produce an output image with the desired colour appearance when the image is viewed on the reference display. The exact imaging chain used to implement such steps can vary widely in complexity and quality. However, in one simple implementation, an automatic white-balance algorithm can be used to estimate channel-dependent gain factors that can be applied to the individual RGB colour channels of the linear sensor RGB values. At this point, neutral scene elements will have approximately equal colour values. Next a sensor characterization matrix can be applied to the linear sensor RGB colour values to determine approximate scene tristimulus values. (Such a matrix can be determined for a nominal scene illuminant using methods such as those being developed for ISO 17321-2.) The scene tristimulus values can next be transformed to the RGB primaries of the softcopy display using a phosphor matrix associated with the RGB primaries of the softcopy display. (Typically, the sensor characterization matrix and the phosphor matrix can be combined into a single matrix operation for implementation purposes.) Next, a simple colour-rendering transformation can be implemented using a one-dimensional look-up table (LUT) that applies a tone scale function to map the scene radiance values to corresponding desired image radiance values. Finally, an appropriate non-linear display transfer function (such as that specified for an sRGB display) can be applied to account for the non-linear response of the softcopy display. (The tone scale function and the non-linear display transfer function can be combined into a single one-dimensional look-up table for implementation purposes.) This entire sequence of transformation can therefore be implemented using a simple LUT-matrix-LUT imaging chain.

A.4 Generic workflow for copy-stand photography

Figure A.3 shows a workflow diagram for using a digital camera in a copy-stand photography application. While the workflow is quite similar to that shown in Figure A.1, there are several key differences. The most important difference is that the input image is a piece of original artwork (e.g. a photograph, a painting or a drawing) rather than a scene. As a result, the input image is already in a picture-referred image state, and therefore it is inappropriate to apply a colour-rendering operation. In this case, the original picture has presumably already been rendered so that it has the desired colour appearance (at least under the appropriate original viewing conditions). Depending on the characteristics of the input and the output media, and any differences in the original and output-viewing conditions, a colour re-rendering transform may be needed to determine an image appropriate for the output medium and viewing conditions.

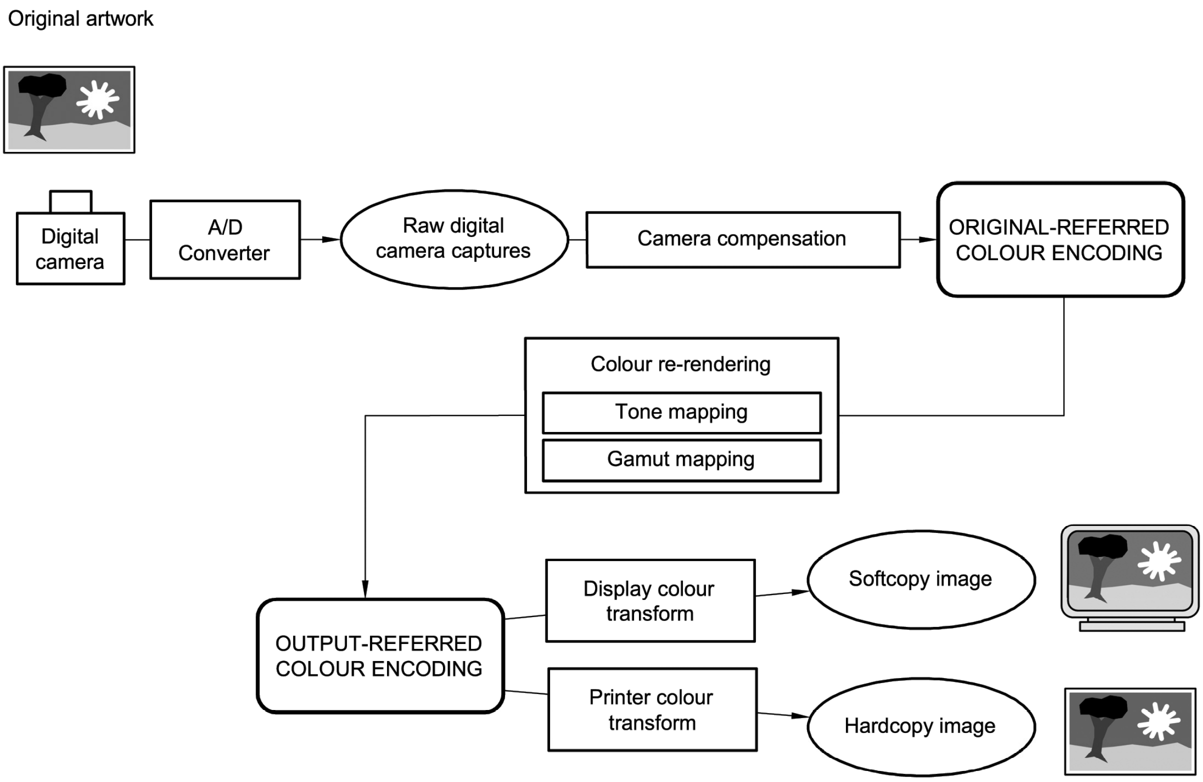

Figure A.3 — Generic workflow for copy-stand photography

Since copy-stand photography applications are usually carried out under well-controlled image capture conditions, the camera compensation transform can often be determined for the actual conditions under which the camera will be used. As a result, automatic exposure control and white balance operations will typically be unnecessary, and the accuracy of the camera compensation transform may be significantly higher. However, unless the digital camera has spectral sensitivities that are linear combinations of CIE colour-matching functions, the optimal camera compensation transform may be input-media-dependent.

The following should be noted.

— In many cases, the colour re-rendering transform may be close to an identity transform if the input and output media and viewing conditions are sufficiently similar.

— Since the goal of most copy-stand photography applications is to produce an exact copy of the original artwork, there is usually no need for image-editing operations, or colour preference transformations, although such steps can be useful to correct defective originals or make desirable image enhancements.

A.5 Generic workflow for scanning hardcopy

Figure A.4 shows an analogous workflow diagram for an application using a hardcopy scanner to scan an original image. This workflow is almost identical to that shown in Figure A.3 for copy-stand photography, except that the digital camera has been replaced by a hardcopy scanner, and the camera compensation transform has been replaced with a corresponding scanner compensation transform.

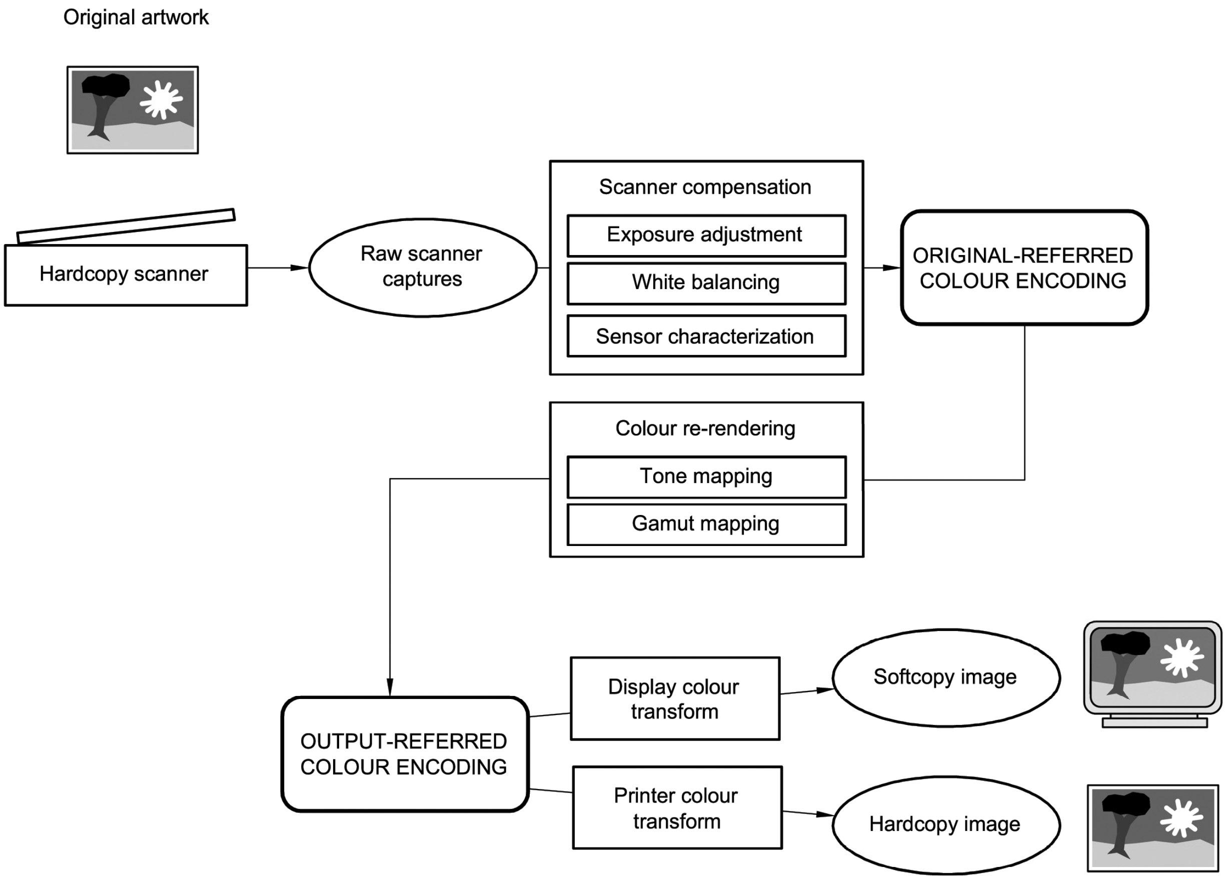

Figure A.4 — Generic workflow for scanning hardcopy

It should be noted that the accuracy of the scanner compensation transformation will be a function of the scanner spectral sensitivities and the input media characteristics, as well as the accuracy and robustness of the modelling technique used to determine the scanner compensation transformation. Consequently, the accuracy of the original-referred image data will also depend on these same factors. The accuracy requirements for the original-referred image data will be application-dependent.

A transparency scanner workflow would be very similar to the hardcopy scanner workflow, except that the colour re-rendering transform would typically need to make more significant changes since the dynamic range and intended viewing conditions for transparency materials are typically quite different than that of most other types of output media.

A.6 Generic workflow for hardcopy scanners producing display-ready images

Depending on the application, other workflows may also be appropriate for use with hardcopy scanners. For example, Figure A.5 shows a workflow that is common in many desktop scanner applications. In this case, the scanner produces a finished file in an output-referred display-ready colour encoding (e.g. sRGB). The in-scanner processing converts the raw scanner captures all the way through to the final display-ready colour encoding. The in-camera processing will typically include any required scanner compensation and colour re-rendering operations to produce an output image with the desired colour appearance when the image is viewed on the reference display. (The in-scanner processing may be applied in the scanner itself, or in accompanying host software.). Since images produced by the scanner are in a display-ready colour encoding, they can be displayed directly on a softcopy display without the need for additional colour transforms (provided that the characteristics of the softcopy display are sufficiently close to those of the standard display assumed for the display-ready colour encoding). A printer colour transform is still required to print hardcopy images, but in most desktop digital imaging applications, this colour transform is applied in the printer driver or hardware and is invisible to the user.

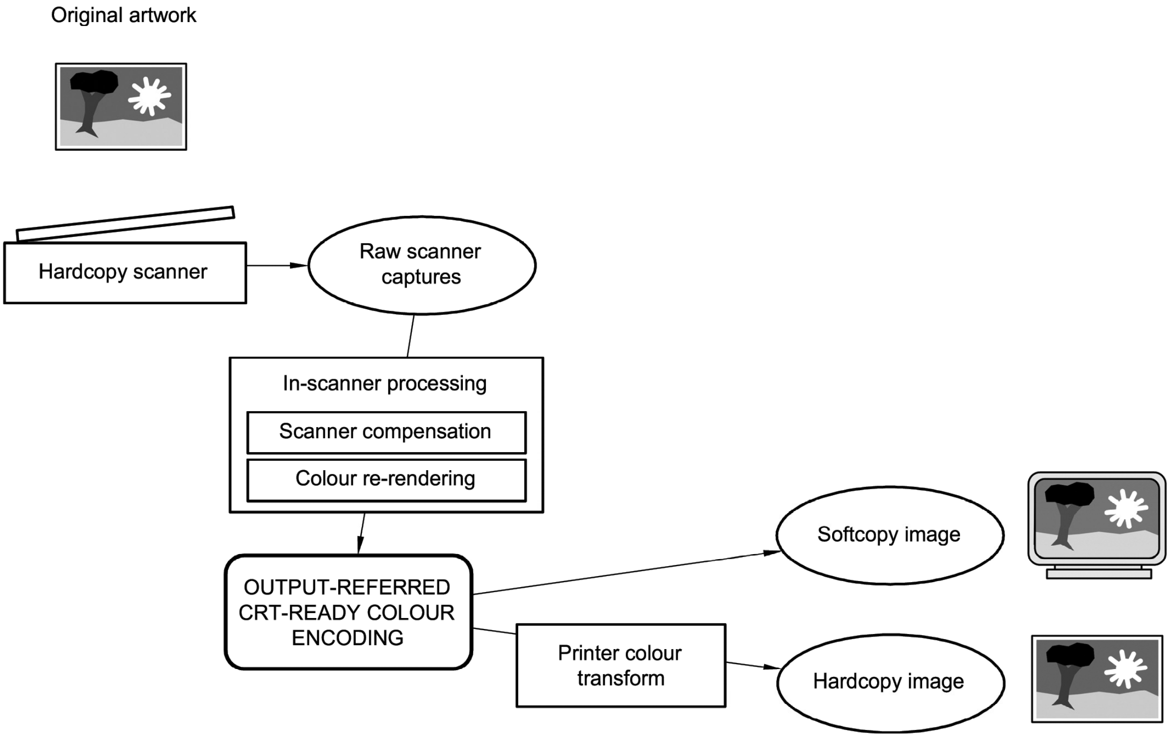

Figure A.5 — Generic workflow for hardcopy scanners producing display-ready images

A.7 Generic workflow for scanning colour negatives

Figure A.6 shows a generic workflow diagram for scanning colour negatives. This workflow has some notable differences relative to the other workflows that have been discussed due to the fundamentally different characteristics of colour negatives. The diagram actually illustrates two different workflows:

— the first going through a scene-referred image state;

— the second going directly to an output-referred image state.

Due to the extremely wide dynamic range of colour negative film, it represents an excellent means for capturing scene-referred image data. To back out the characteristics of the film and scanner, a film unrendering transform is needed to infer the original scene colours from the film scan. In addition to compensating for the colour and tone reproduction characteristics of the film, it is usually necessary to account for the optical flare in the film camera and the colour response of the scanner. Also, it is usually necessary to apply a colour/density balancing step to account for the fact that colour negatives are typically captured over a wide range of exposure levels and using a wide variety of illuminants. It should be noted that the workflow diagram shows scanner compensation, colour/density balancing and film compensation as being components of the film unrendering transform. However, they could also be implemented as a series of individual transforms.

For the second workflow shown in Figure A.6, the film scan is used to directly form output-referred image data. In this case, a film rendering transform is needed to determine the desired picture colorimetry from the film scan. Typically, such transforms include a print model to mimic the response of photographic paper exposed using a conventional photographic enlarger. As with the film unrendering path, it is also necessary to apply scanner compensation, and colour/density balancing operations.

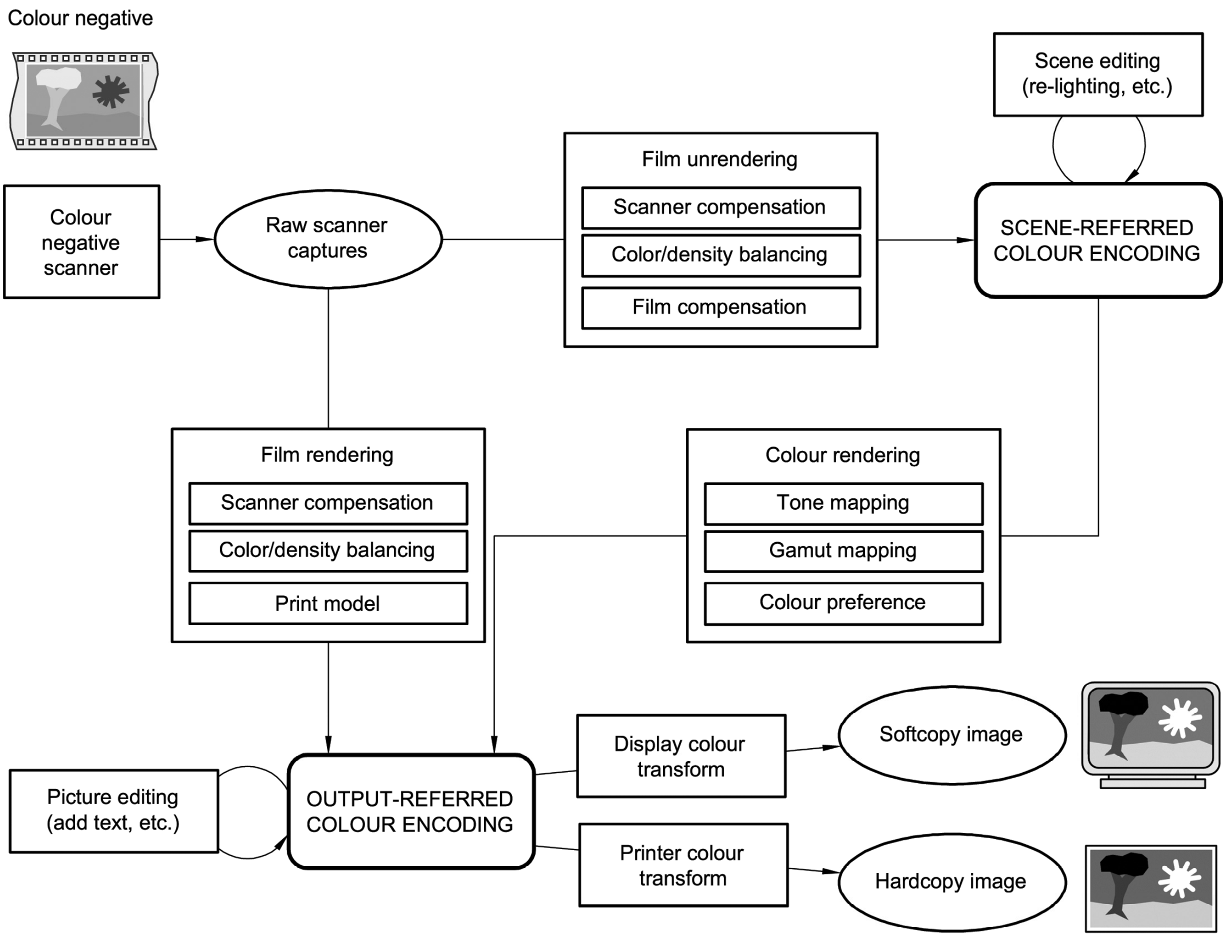

Figure A.6 — Generic workflow for scanning colour negatives.

The following should be noted.

— The film unrendering workflow maintains the greatest amount of flexibility for repurposing and editing the image data. However, it also requires a more careful characterization of the film response in order to create the film unrendering transform.

— For cases where it is a requirement to match conventional photographic prints, the film rendering path will generally be a more straightforward solution.

A.8 Generic workflow for video imaging systems

Figure A.7 shows a workflow diagram for video imaging systems. While, in many ways, video imaging systems are closely related to the digital camera systems shown in Figures A.1 and A.2, there are several points that are worth further discussion due to the potential for confusion when fitting such systems into the image-state-based digital imaging architecture shown in Figure 1. Most video system standards, such as Rec. ITU-R BT.709, are actually video camera standards, and as such specify the relationship between the colour of a scene and the colour encoding that video cameras should produce. Therefore, according to the guidelines presented in this part of ISO 22028, such video RGB colour encodings should strictly be classified as “scene-referred” colour encodings. However, since these same video RGB colour encodings are adapted to be ready for direct display on a standard television, it could also be argued that they are “output-referred” colour encodings as well. This has led some to suggest that the distinction between scene-referred images and output-referred images is artificial (and perhaps unnecessary), and additionally that the image-state-based digital imaging architecture does not adequately represent some important classes of imaging systems.

However, it is important to remember that while most colour encodings are intended to be used to encode images in a particular image state, it is not the colour encoding itself that is in a particular image state, but rather the colorimetry that the colour encoding is representing. As was discussed in 5.4.3, it is possible to specify the relationship between a colour encoding and the image colorimetry of the image in multiple image states. This is effectively the situation in video imaging systems. The relationship between the video RGB values and the original scene colorimetry is specified by video camera standards, such as Rec. ITU-R BT.709, while the relationship between the video RGB values and the resulting display colorimetry is specified by the colour response characteristics of a typical display. (Display colour response characteristics are largely non-standardized. However, the sRGB colour-encoding standard, IEC 61966-2-1, is intended to be representative of the response of a typical computer-controlled CRT.) Therefore, it can be argued that video RGB colour encodings are both scene-referred and output-referred.

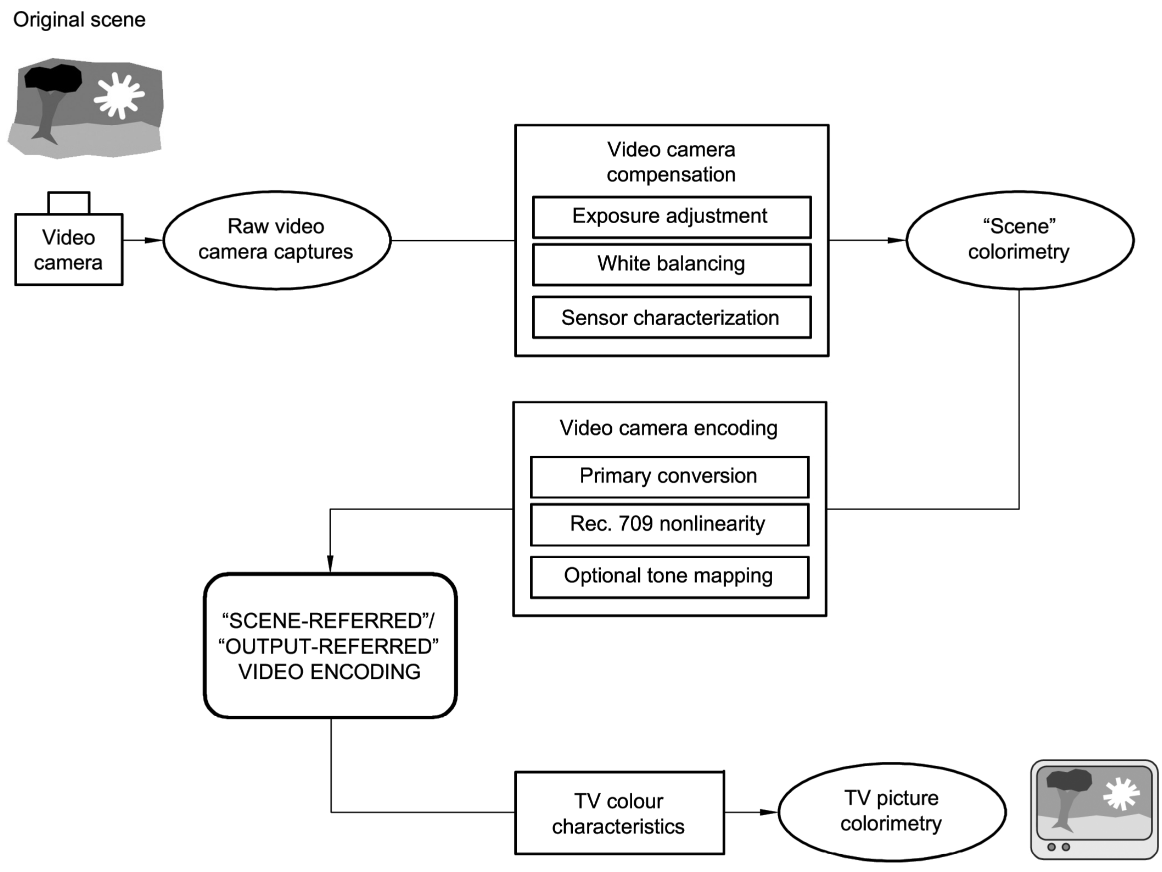

Figure A.7 — Generic workflow for video imaging systems.

However, even if it is accepted that video RGB colour encodings can be both scene-referred and output-referred, there is another point of potential confusion that frequently arises. Since the same colour encoding effectively is tied to multiple image states, this has led some to suggest that there is no colour-rendering step being applied in video imaging systems, and therefore that good image quality can be achieved even in the absence of a colour-rendering step. However, it is important to remember that even though there may be an identity transformation between the “video camera RGB values” and the “video display RGB values,” this does not imply that there is an identity transformation between the original scene colorimetry and the display colorimetry. This is because the video camera encoding characteristics (specifying the relationship between the video RGB values and the scene colorimetry) are generally not the inverse of the display characteristics (specifying the relationship between the video RGB values and the display picture colorimetry).

Consider the workflow diagram shown in Figure A.7. A video camera is used to capture an image of an original scene. Generally some type of video camera compensation step is used to process the raw video camera sensor signals, accounting for the spectral response of the particular sensors, and performing operations such as exposure compensation and white balance. The result of these operations is an estimate of the “scene” colorimetry.

It should be noted that the “scene” colorimetry in this case may not be an accurate representation of the actual colorimetry for the original scene. For example, the estimated “scene” colorimetry may be relative to a different adopted white point than that of the actual scene. Additionally, there may be inaccuracies in the estimated scene colorimetry due to factors such as uncorrected optical blurring and flare, noise, and colour analysis errors due to capture device metamerism.

The scene colorimetry is then transformed to a standard video RGB colour encoding according to an appropriate colour-encoding standard such as Rec. ITU-R BT.709. This transformation generally includes a primary conversion step, together with a nonlinear “gamma correction” function. Additionally, many high-end video systems may include additional tone mapping to smoothly compress the highlight information rather than harshly clipping the colour values.

In many video camera processing chains, there may not be a clear distinction between the video camera compensation and video camera encoding steps. In fact, any particular operation in the image-processing chain may achieve elements of both steps. For example, a single “colour correction matrix” may be used to compensate for the sensor spectral response and convert to the appropriate set of RGB colour space primaries.

Since the resulting “scene-referred” video camera RGB colour values are designed to be ready for direct input to a display, “output-referred” display RGB colour values are conceptually determined by applying an identity transform. However, as discussed earlier, even though an identity transform is used to convert between the scene-referred to the output-referred colour encodings, this does not imply an identity colour-rendering transform. Indeed, since the television display characteristics are generally not designed to be the inverse of the video camera encoding characteristics, the resulting television picture colorimetry will not be equal to the original scene colorimetry. For example, combining the nonlinear encoding function associated with the Rec. ITU-R BT.709 video camera standard, together with a typical nonlinear display transfer function (such as that specified for an sRGB display) results in an effective colour-rendering transform with an overall boost in luminance and chrominance contrast similar to that produced by most other conventional imaging systems. (For a more detailed discussion, see Reference [13].)

Therefore, it can be seen that while it is not possible to point to a particular “colour-rendering step” in the video imaging chain, this does not change the fact that colour-rendering is still occurring. It is the co-optimized design of the video camera encoding and the television display characteristics that produces the desired colour-rendering characteristics. Effectively, this spreads the responsibility for colour-rendering across two components of the imaging system. However, in practice, it is the video camera processing that bears the brunt of the rendering duties. While the theory would dictate that the video RGB colour values produced by a video camera have a defined relationship to the colorimetry of the scene, in reality, the fundamental responsibility of the video camera is to produce a pleasing result when the image is displayed on a typical television. As a result, high quality video cameras will often deviate from the theoretical encoding relationship to incorporate many of the desirable characteristics of typical colour-rendering transforms, such as smooth highlight compression.

Therefore, while it can be justified to label video RGB as both a scene-referred colour encoding and an output-referred colour encoding, in practice video RGB colour encodings have much more in common with output-referred colour encodings than they do with scene-referred colour encodings. They also have almost none of the advantages typically associated with scene-referred colour encodings, such as extended dynamic range and colour gamut, and the flexibility to delay final colour-rendering decisions until a later time. As a result, it is not possible to obtain the full benefits associated with a scene-referred colour encoding when video RGB colour encodings are used for interchange of images in an open systems environment.

The key point to understand is that it is not the colour encoding itself that is in one image state or another, but rather the underlying image colorimetry that the colour encoding is representing. Therefore, while some colour encodings, such as video RGB, can be linked to multiple image states, the image colorimetry obtained by applying a specified colour transformation to the encoded colour values should be associated with a single image state.

Annex B
(informative)

Characteristics of existing colour encodings

A number of colour encodings have been defined by various standards organizations and industry consortia. To enable a comparison of these colour encodings, Tables B.1 to B.5 summarize the characteristics of each colour encoding relative to the terminology and requirements defined in this part of ISO 22028.

The following should be noted.

— In some cases where a particular characteristic was not explicitly specified, an appropriate value could be inferred from other information that was provided. In such cases, a footnote is included to indicate how the quantity in the table was derived.

— The specifications for these colour encodings vary in their degree of completeness relative to the requirements set forth in this part of ISO 22028 for defining an unambiguous colour image encoding. Therefore, in the terminology of this document, some of these colour encodings should be classified as colour image encodings and others as colour space encodings.

Table B.1 — Characteristics of sRGB, sYCC and bg-sRGB colour encodings

Attribute	sRGB	sYCC	bg-sRGB
Standard	IEC 61966-2-1	IEC 61966-2-1, Amendment 1 (similar to sRGB YCC in PIMA 7667:2001)	IEC 61966-2-1, Amendment 1 (similar to e-sRB in PIMA 7667:2001)
Colour space characteristics			
Colour space type	Colorimetric: RGB colour space	Colorimetric: Luma-chroma colour space	Colorimetric: RGB colour space
Colour gamut	Video display-based	Extended	
RGB primaries	R: $x =$ 0,640 0, $y =$ 0,330 0 G: $x =$ 0,300 0, $y =$ 0,600 0 B: $x =$ 0,150 0, $y =$ 0,060 0 (from ITU-R BT.709-3)		
Colour component transfer function	$C =$ 12,92 $\times C$ for $C \leqslant$ 0,003 130 8 $C =$ 1,055 $\times C^{1/2,4} -$ 0,055 for $C >$ 0,003 130 8	$C = -$1,055 $\times$ ($- C)^{1/2,4} +$ 0,055; for $C < -$0,003 130 8 $C =$ 12,92 $\times C$; for $\lvert C \rvert \leqslant$ 0,003 130 8 $C =$ 1,055 $\times C^{1/2,4} -$ 0,055; for $C >$ 0,003 130 8 *(extended from sRGB)*	
Luma-chroma matrix	N/A	$Y =$ 0,299 0R' $+$0,587G' $+$ 0,114 0B' $C_R =$ 0,500 0R' $-$ 0,418 7G' $-$0,081 3B' $C_B = -$0,168 7R' $-$ 0,331 2G' $+$0,500 0B' (from ITU-R BT601-5)	N/A
Colour space white point luminance	80 cd/m^3		
Colour space white point chromaticity	$x =$ 0,312 7, $y =$ 0,329 0 (D_{65})		
Colour space encoding characteristics			
Colour space value range	Linear RGB: 0,0 to 1,0	Linear Y: 0,0 to 1,0	Linear RGB: $-$0,53 to 1,68
Digital code value range	0 to 255 (8-bit) provided as an example others allowed with WDC $= 2^n - 1$ and KDC = 0		0 to 1 023 (10-bit) (12-bit and 16-bit versions also given in PIMA 767:2001)
Colour image encoding characteristics			
Image state	Output-referred (CRT)	Output-referred (hypothetical extended gamut CRT)	Output-referred (hypothetical extended gamut CRT)
Image surround	*“Background”* $-$20 % of display white point luminance level (16 cd/m^2) *“Surround”* $-$20 % reflectance of ambient luminance level (4,1 cd/m^2)		
Adapted white point luminance [a]	Unspecified [b]		
Adapted white point chromaticity [a]	Unspecified [b]		
Luminance of adapting field	Unspecified [b]		
Viewing flare (typical viewing conditions)	6,9 % of colour space white point luminance [c]		
Valid relative luminance range (excluding viewing flare and veiling glare)	0,0 to 1,0		
Reference medium white point luminance	80 cd/m^2		
Reference medium white point chromaticity	$x =$ 0,312 7, $y =$ 0,329 0 (D_{65})		
Reference medium black point luminance	1,0 cd/m^2 [d]		
Reference medium black point chromaticity	$x =$ 0,312 7, $y =$ 0,329 0 (D_{65}) [d]		

Table B.1 — Characteristics of sRGB, sYCC and bg-sRGB colour encodings (*Continued*)

[a] The image adapted white point is the stimulus that an observer who is adapted to the viewing environment, where the image is viewed, would judge to be neutral and have a luminance factor of unity; i.e. would be perceived to be a perfect white diffuser. Note that the image adapted white will not always correspond to actual measurements of a perfect reflecting diffuser in the image viewing environment. Also note that the mapping of the adapted white for a scene to particular colorimetric values in an output-referred image is a part of the proprietary colour-rendering. It will frequently be undesirable to map the scene adapted white point directly to the image adapted white point.

[b] The sRGB, sYCC, and bg-sRGB colour encodings do not explicitly identify an adapted white point. It can probably be inferred that the adapted white point chromaticity would be equal to that of the media white point, although the adapted white point luminance is open to interpretation. The adapted white point would be influenced by the media white point and the standard viewing conditions, and additionally by other factors such as the image presentation. Likewise, the luminance of the adapting field is also not specified, but will be closely linked to the assumed adapted white point.

[c] This value was determined by dividing the specified veiling glare of 5,5 cd/m^2 by the specified encoding white point luminance of 80 cd/m^2.

[d] This reference medium black point corresponds to the black point of a typical CRT display in the sRGB reference viewing conditions (with a D_{65} chromaticity to maintain the neutral axis), as measured using a tele-spectral-radiometer from the observer position. This value is not specified in IEC 61966-2-1, and is called the viewer observed black point in PIMA 7667. See Annex D of PIMA 7667 for more information.

Table B.2 — Characteristics of scRGB, scRGB-nl and sYCC-nl colour encodings

Attribute	sRGB	scRGB-nl	scYCC-nl
Standard	IEC 61966-2-2		
Colour space characteristics			
Colour space type	Colorimetric: RGB colour space		Colorimetric: Luma-chroma colour space
Colour gamut	Extended		
RGB primaries	R: $x = 0{,}640\ 0$, $y = 0{,}330\ 0$ G: $x = 0{,}300\ 0$, $y = 0{,}600\ 0$ B: $x = 0{,}150\ 0$, $y = 0{,}060\ 0$ (from ITU-R BT.709-3)		
Colour component transfer function	Linear	$C = -1{,}055 \times (-C)^{1/2{,}4} + 0{,}055$; for $C < -0{,}003\ 130\ 8$ $C = 12{,}92 \times C$; for $\lvert C \rvert \leqslant 0{,}003\ 130\ 8$ $C = 1{,}055 \times C^{1/2{,}4} - 0{,}055$; for $C > 0{,}003\ 130\ 8$ *(extended from sRGB)*	
Luma-chroma matrix	N/A		$Y = 0{,}299\ 0R'$ $+0{,}587\ 0G' + 0{,}114\ 0B'$ $C_R = 0{,}500\ 0R' - 0{,}418\ 7G'$ $-0{,}081\ 3B'$ $C_B = -0{,}168\ 7R' - 0{,}331\ 2G'$ $+0{,}500\ 0B'$ (from ITU-R BT601-5)
Colour space white point luminance	Absolute white point luminance values are scene-dependent		
Colour space white point chromaticity	$x = 0{,}312\ 7$, $y = 0{,}329\ 0$ (D_{65}) [a]		
Colour space encoding characteristics			
Colour space value range	Linear RGB: $-0{,}5$ to 7,499 9	Linear RGB: $-0{,}603\ 8$ to 7,591 3	Linear Y: $-603\ 8$ to 7,591 3
Digital code value range	0 to 65 535 (16-bit)		0 to 4 095 (12-bit)
Colour image encoding characteristics			
Image state	Scene-referred		
Image surround	Unspecified		
Adapted white point luminance	Unspecified [b]		
Adapted white point chromaticity	Unspecified [b]		
Luminance of adapting field	Unspecified [b]		
Viewing flare	N/A		
Valid relative luminance range	0,0 to 7,499 9 [c]	0,0 to 7,591 3 [c]	
Reference medium white point luminance	N/A		
Reference medium white point chromaticity	N/A		
Reference medium black point luminance	N/A		
Reference medium black point chromaticity	N/A		

[a] The scRGB, scRGB-nl, and scYCC-nl colour encodings do not explicitly identify colour space white point chromaticity values; however, it can be inferred from the XYZ-to-RGB matrix that they correspond to D_{65}.

[b] The scRGB, scRGB-nl, and scYCC-nl colour encodings do not explicitly identify an adapted white point, because the adapted white point can vary within a scene, and methods for the determination of the adapted white point are proprietary. Ideally, the adapted white point would be equal to the colour space white point. In practice, the adapted white point results from the camera white balance and exposure settings, and is open to interpretation. Likewise, the luminance of the adapting field is also not specified, but will be closely linked to the assumed adapted white point.

[c] The scRGB, scRGB-nl, and scYCC-nl colour encoding specifications provide for the encoding of negative RGB values. However, since colours with negative luminance values are physically impossible, such colour values will not occur in scenes and therefore should not be used in practice, so the valid luminance range would be limited between 0,0 and the top of the corresponding encoding range.

Table B.3 — Characteristics of ROMM RGB, RIMM RGB and ERIMM RGB colour encodings

Attribute	ROMM RGB	RIMM RGB	ERIMM RGB
Standard	ANSI/I3A IT10.7666	ANSI/I3A IT10.7466	
Colour space characteristics			
Colour space type	Colorimetric: RGB colour space		
Colour gamut	Extended		
RGB primaries	R: $x = 0{,}734\,7$, $y = 0{,}265\,3$ G: $x = 0{,}159\,6$, $y = 0{,}840\,4$ B: $x = 0{,}036\,6$, $y = 0{,}000\,1$		
Colour component transfer function	$C = 16 \times C$ for $C < 0{,}001\,953$ $C = C^{1/1{,}8}$ for $C \geqslant 0{,}001\,953$	$C = (45 \times C)/1{,}402$ for $C \leqslant 0{,}018$ $C = (1{,}099 \times C^{0{,}45} - 0{,}099)/1{,}402$ for $C > 0{,}018$ (from ITU-R BT.709-3)	$C = 29{,}048\,7 \times C$ for $C \leqslant 0{,}002\,718\,28$ $C = (\log_{10} C + 3)/5{,}5$ for $C > 0{,}002\,718\,28$
Luma-chroma matrix	N/A		
Colour space white point luminance	142 cd/m^2 [a]	15 000 cd/m^2	
Colour space white point chromaticity	$x = 0{,}345\,7$, $y = 0{,}358\,5$ (D_{50})		
Colour space encoding characteristics			
Colour space value range	Linear RGB: 0,0 to 1,0	Linear RGB: 0,0 to 2,0	Linear RGB: 0,0 to 312,2
Digital code value range	0 to 255 (8-bit); 0 to 4 095 (12-bit); 0 to 65 535 (16-bit)		0 to 4 095 (12-bit); 0 to 65 535 (16-bit)
Colour image encoding characteristics			
Image state	Output-referred (print)	Scene-referred	
Image surround	“Average” −20 % of the adapted white point		
Adapted white point luminance	160 cd/m^2	15 000 cd/m^2	
Adapted white point chromaticity	$x = 0{,}345\,7$, $y = 0{,}358\,5$ (D_{50})		
Luminance of adapting field	Not explicitly specified, but a default level of 20 % of the adapted white point luminance can be assumed in most cases		
Viewing flare	0,75 % of adapted white point	N/A	
Valid relative luminance range (excluding viewing flare and veiling glare)	0,0 to 1,0	0,0 to 2,0	0,0 to 316,2
Reference medium white point luminance	142 cd/m^2 [a]	N/A	
Reference medium white point chromaticity	$x = 0{,}345\,7$, $y = 0{,}358\,0$ (D_{50})	N/A	
Reference medium black point luminance	0,5 cd/m^2 [b]	N/A	
Reference medium black point chromaticity	$x = 0{,}345\,7$, $y = 0{,}358\,5$ (D_{50})	N/A	

[a] Equal to the luminance of the adapted white times the luminance factor or the reference medium white point.

[b] Equal to the luminance of the adapted white times the luminance factor or the reference medium black point.

Table B.4 — Characteristics of ICC PCS (v4) colour encodings

Attribute	ICC v4 Perceptual Intent PCS [a]	ICC v4 Colorimetric Intent PCS [a]
Standard	ICC.1:2001	
Colour space characteristics		
Colour space type	Colorimetric: Cie colour space (linear XYZ or L*a*b*)	
Colour gamut	Extended	
RGB primaries	N/A	
Colour component transfer function	N/A	
Luma-chroma matrix	N/A	
Colour space white point luminance	142 cd/m^2	Variable [f]
Colour space white point chromaticity	$x =$ 0,345 7, $y =$ 0,358 5 (D_{50})	Variable [f]
Colour space encoding characteristics		
Colour space value range	Linear Y: 0,0 to 1,0 [e]	
Digital code value range	Variable	
Colour image encoding characteristics		
Image state	Output-referred (print)	Original-referred
Image surround	"Average" $-$20 % of the adapted white point	Variable [f]
Adapted white point luminance	160 cd/m^2	Unspecified [b]
Adapted white point chromaticity	$x =$ 0,345 7, $y =$ 0,358 5 (D_{50})	Unspecified [b]
Luminance of adapting field	Not explicitly specified, but a level of 20 % of the adapted white point luminance can be assumed	Unspecified [b]
Viewing flare	0,75 % of adapted white point	Unspecified
Valid relative luminance range (excluding viewing flare and veiling glare)	0,0 to 1,0	
Reference medium white point luminance	142 cd/m^2 [c]	N/A [g]
Reference medium white point chromaticity	$x =$ 0,345 7, $y =$ 0,358 5 (D_{50})	N/A [g]
Reference medium black point luminance	0,5 cd/m^2 [d]	N/A [g]
Reference medium black point chromaticity	$x =$ 0,345 7, $y =$ 0,358 5 (D_{50})	N/A [g]

[a] ICC PCS is generally not used for the encoding of digital images, but it serves a similar purpose in its role of being a connection point for colour transforms.

[b] The ICC v4 Colorimetric Intent PCS colour encoding does not explicitly identify an adapted white point. Proprietary methods could be used to estimate the adapted white point from the "viewingConditionsTag", the "mediaWhitePointTag" and the image data. Likewise, the luminance of the adapting field is also not specified, but will be closely linked to the estimated adapted white point.

[c] Equal to the luminance of the adapted white times the luminance factor of the reference medium white point.

[d] Equal to the luminance of the adapted white times the luminance factor of the reference medium black point.

[e] Where $Y =$ 1,0 corresponds to the medium white point luminance.

[f] The colour space white point and the image surround for the ICC v4 Colorimetric Intent PCS colour encoding are variable and depend on "mediaWhitePointTag" and "viewingConditionsTag".

[g] The ICC v4 Colorimetric Intent PCS colour encoding does not have a defined reference medium. The actual medium white point and black point will be variable and will depend on the ICC "mediaWhitePointTag" and ICC "mediaBlackPointTag."

Table B.5 — Characteristics of ITU-R BT.709 and ITU-R BT.601 colour encodings

Attribute	ITU-R BT.709 RGB	ITU-R BT.709 YC_RC_B	ITU-R BT.601 YC_RC_B [a]
Standard	ITU-R BT.709-3		ITU-R BT.601-5
Colour space characteristics			
Colour space type	Colorimetric: RGB colour space	Colorimetric: Luma-chroma colour space	
Colour gamut	video display-based [b]		
RGB primaries	R: $x = 0{,}640$, $y = 0{,}330$ G: $x = 0{,}300$, $y = 0{,}600$ B: $x = 0{,}150$, $y = 0{,}060$		Unspecified
Colour component transfer function [c]	$C' = 4{,}500 \times C$; for $0 \leqslant C < 0{,}018$ $C' = 1{,}099 \times C^{0{,}45} - 0{,}099$; for $0{,}018 \leqslant C \leqslant 1$		Unspecified
Luma-chroma matrix	N/A	$Y = 0{,}212\,5R' + 0{,}715\,2G' + 0{,}072\,2B'$ $C_R = -0{,}635\,0(R' - Y)$ $C_B = 0{,}538\,9(B' - Y)$	$Y = 0{,}299R' + 0{,}587G' + 0{,}114B'$ $C_R = 0{,}500R' - 0{,}419G' - 0{,}081B'$ $C_B = -0{,}169R' - 0{,}331G' + 0{,}500B'$
Colour space white point luminance	Unspecified		
Colour space white point chromaticity	$x = 0{,}312\,7$, $y = 0{,}329\,0$ (D_{65})		Unspecified
Colour space encoding characteristics			
Colour space value range	Not clearly specified — Nominally linear RGB: 0,0 to 1,0	Not clearly specified — Nominally linear Y: 0,0 to 1,0	
Digital code value range	1 to 254 (8-bit) ("Black level" = 16 "Nominal peak" = 235) [d] 4 to 1 019 (10-bit) ("Black level" = 64 "Nominal peak" = 940) [d]	1 to 254 (8-bit) ("Black level" for $Y = 16$ "Nominal peak" for $Y = 235$ "Nominal peak" for C_R and $C_B = 16$ and 240) [d] 4 to 1 019 (10-bit) ("Black level" for $Y = 64$ "Nominal peak" for $Y = 940$ "Nominal peak" for C_R and $C_B = 64$ and 960) [d]	1 to 254 (8-bit) ("Black level" for $Y = 16$ "Nominal peak" for $Y = 235$ "Nominal peak" for C_R and $C_B = 16$ and 240) [d]
Colour image encoding characteristics			
Image state	Scene-referred [e]		
Image surround	Unspecified		
Adapted white point luminance	Unspecified		
Adapted white point chromaticity	Unspecified		
Luminance of adapting field	Unspecified		
Viewing flare	N/A		
Valid relative luminance range	Unspecified		
Reference medium white point luminance	N/A		
Reference medium white point chromaticity	N/A		
Reference medium black point luminance	N/A		
Reference medium black point chromaticity	N/A		

[a] According to the terminology used in this part of ISO 22028, ITU-R BT.601 YC_RC_B is neither a colour space nor a colour encoding since it is not based upon a defined colour space. Rather it is a luma-chroma transformation and a digital encoding method that can be applied to a video RGB colour space such as ITU-R BT.709 RGB to determine a corresponding luma-chroma colour space encoding.

[b] The luma-chroma colour encodings permit the encoding of an extended colour gamut. However, ITU-R BT.601-5 recommends that the gamut be limited to that of the RGB representation.

[c] The colour component transfer function built into many real video capture systems may differ somewhat from those specified in the standards in order to achieve a desired tone mapping on typical video displays. For example, some colour component transfer functions incorporate features such as smooth roll-off of highlight detail.

[d] The standard specifies digital encoding values for the "black level" and the "nominal peak." However, it is unclear whether these points should be interpreted to correspond to a colour space value range of 0,0 to 1,0 or to some other colour space value range associated with real/ideal stimuli. Digital encoding values outside of the "black level" to "peak level" range are apparently permitted, presumably corresponding to extended range scene data.

[e] Video colour encoding can also be viewed as output-referred since they are designed to be compatible with typical video displays. In many cases, this may be a more accurate characterization, particularly if the video capture system has used non-standard encoding functions to produce a more pleasing image on typical video displays. For more information, see discussion in A.8.

Annex C
(informative)

Criteria for selection of colour encoding

C.1 General

A set of guiding principles for making the appropriate selection of colour encodings for image storage, manipulation and interchange in digital photography and graphic technology applications is as follows.

— Retain as much colour information as needed for as long as needed.

— Facilitate the compositing of images from different sources.

— Preserve the option of going out to a multiplicity of output devices.

— Communicate colour appearance, not just colorimetry.

— Facilitate the editing of images so that image quality is maximized.

— Minimize computational and storage overhead.

— Maximize interoperability across vendors, applications and products.

— Ensure compatibility with standards and other common workflows.

When identifying colour encodings for use in a particular workflow, a number of different selection criteria should be considered. The importance of any one criterion will be a function of the application. The following sections briefly describe some of the selection criteria that will be important in many digital photography and graphic technology applications.

It should be noted that CIE TC8-05 is in the process of developing and documenting quantitative metrics that can be used for the evaluation and comparison of various colour encodings. The metrics that they are compiling should be directly applicable for the evaluation of many of the criteria that will be discussed in the following sections.

C.2 Image state

The image state associated with the colour image encoding will be an important selection criterion for many applications. The workflow for a particular application may revolve around images in one particular image state, or alternatively may pass through several image states where appropriate operations are applied at each stage.

In some workflows an important objective will be to capture and preserve as much information about an original scene or an original piece of artwork as possible. This goal is best satisfied by a scene-referred colour encoding for original scenes, or an original-referred colour encodings for original artwork. By deferring decisions about colour-rendering (or re-rendering), flexibility for optimally using images for a variety of purposes can be preserved. Images stored in scene-referred or original-referred colour encodings can be archived for later use without needing to commit to what uses they will be put.

Another important objective in some workflows is to define a colour image encoding for unambiguous representation of the intended colour appearance of output-referred pictures. In this case, all colour-rendering and re-rendering has been applied to produce the desired output picture on the reference medium. The choice of the reference medium affects the image's ability to be used for different purposes. A reference medium similar to some commonly occurring real output medium can facilitate simple colour re-rendering to the real medium. On the other hand, a reference medium with a large (but still realistic) dynamic range and gamut may be capable of producing somewhat more accurate colour reproduction over a wide range of real output media, although it may require more complex colour re-rendering.

C.3 Extent of colour gamut

Digital photography and graphic technology applications employ devices and materials having a diverse set of colour gamuts. For example, the colour gamuts of various softcopy and hardcopy output devices differ markedly in both size and shape. Additionally, the gamut of colours that occur in real scenes is even more extensive due to the wide variety of reflective, transmissive, and emissive surfaces that can be encountered, as well as the wide range of lighting conditions. The extent of the colour gamut associated with a colour encoding is therefore a very important criterion that should be considered in the selection of a colour encoding for a particular application.

The particular set of colours that will be important will vary as a function of application, and may also vary throughout the workflow for a particular application. It is important to select a colour encoding that has a large enough colour gamut to encode the range of colours that are important within a particular workflow. But it is also desirable to select a colour encoding where the colour gamut is no larger than necessary in order to avoid wasting code values on colours that will never get used in practice. This is necessary to maximize encoding efficiency and minimize quantization errors.

Rather than using a brute force metric, such as the total volume of the colour gamut, to compare colour encodings, it is preferable to utilize metrics that compare the colour gamut of the colour encoding to the gamut of colours that are important for the particular application of interest. One useful metric is to compute the fraction of the colours in the gamut of “important colours” that are within the colour gamut of the colour encoding. (The gamut of “important colours” being the set of colours that is important for a given application. In some cases, this may be the colour gamut of all real world surface colours or all colours within the spectrum locus. In other cases, it may be the colour gamut of a particular device or imaging medium that is central to a particular workflow.) This metric can be computed by determining the volume of the overlap gamut for the colour encoding and the important colours in an approximately uniform colour space for the anticipated viewing condition (e.g. CIELAB), and then dividing by the volume of the gamut of important colours. Ideally, the colour encoding gamut should contain 100 % of the gamut of important colours.

A related metric is to compute the fraction of the code value combinations for a particular colour encoding that are used by colours within the gamut of important colours. While the previous metric is an indication of whether the colour encoding covers the set of important colours, this metric is an indication of how efficiently the colour encoding is used to encode those colours. Ideally, the colour encoding should minimize the percentage of wasted code value combinations that are not used by any important colours, while simultaneously maximizing the number of important colours that can be represented.

C.4 Luminance dynamic range

The luminance dynamic range is one particular dimension of the colour gamut corresponding to the range of luminance values that can be represented. The luminance dynamic range is typically measured by the ratio of the luminance values of the brightest and darkest image content that can be represented in a colour encoding, or alternatively by the logarithm of this ratio. [Sometimes the luminance dynamic range will be given in terms of the number of “photographic stops,” which is the binary logarithm ($\log_2$) of the luminance ratio.]

It should be noted that at the low-luminance end, the darkest image content that can be accurately represented in a colour encoding may be limited by the quantization characteristics of the colour encoding and/or noise in the image source rather than the luminance level associated with the smallest code value. In some colour encodings, the smallest code value has a corresponding luminance value of zero, which would seem to imply an infinite dynamic range. However, in practice, the first few code values should not be included in the calculation of the luminance dynamic range due to the fact that the log luminance difference between consecutive code values is so large.

The luminance dynamic ranges of most picture-referred colour encodings are quite similar, and are tied to the characteristics of the imaging medium. On the other hand, the luminance dynamic range associated with scenes varies quite widely and can be significantly larger than those associated with rendered pictures. As a result, the luminance dynamic range is an important metric for comparing scene-referred colour encodings. Scene luminance ratios for an average scene are about 160:1, but can be as low as 10:1 for some scenes or as high as 50 000:1 for others. Consequently, the world as photographed can far exceed the luminance dynamic range of typical output media used in photography and graphic technology. The colour-rendering process is used to map the luminance dynamic range (and colour gamut) of the scene onto that of a real or virtual output medium. However, if we wish to represent the input scene-referred image itself, we need to employ a colour encoding with an extended luminance dynamic range large enough to represent the full dynamic range of expected scene data. Probably the scene capture means with the largest dynamic range would be a photographic negative, which is able to capture scene exposure ranges exceeding fourteen stops.

C.5 Headroom for brighter-than-white highlights

This metric is closely related to the luminance dynamic range criterion discussed in C.4. It relates to the ability of a colour encoding to represent image values brighter than a “perfect white”. In most cases, this is only a useful metric for scene-referred colour encodings since most common output-referred colour encodings are limited by the white point of a real or virtual reflection-print-like output medium. (The exception to this would be for output-referred colour encodings intended to represent the colour appearance for media such as projected slides, back-lit transparencies, high-quality video and motion pictures, which all have the capability of producing the appearance of brightness levels that are brighter than white.)

The need for representing brighter-than-white highlight information in scene-referred images is much more direct since many scene capture devices, such as digital cameras and colour negative film, are capable of capturing exposure values from stimuli that are significantly brighter than a perfect white diffuser in the principal subject area of the scene. For example, a photograph of a back-lit scene may contain exposure values that are many stops brighter than a properly exposed white object in the foreground.

C.6 Quantization error

Another important characteristic of a colour encoding is the quantization error associated with the discrete nature of the digital colour values. The smallest colour difference that can be represented in a given colour encoding will correspond to a single code value change in the colour value.

In order to evaluate the visual significance of the quantization errors associated with a particular colour encoding, the colour errors should be evaluated using a relatively uniform colour difference metric (e.g. CIELAB ΔE^* or ΔE^*_{94}). For integer digital encodings, one useful metric is to calculate the distribution of colour differences for single code value changes of the various colour channels for a set of test colours distributed within the gamut of important colours. Generally, the resulting colour difference can vary quite widely as a function of location within colour space. Statistics such as the mean and maximum colour differences can then be computed as summary measures of the quantization error. The quantization error will be directly proportional to the number of digital code values, so these metrics will be a strong function of the colour encoding bit-depth.

Besides the bit-depth, the two colour encoding characteristics that will have the most direct impact on the quantization error statistics are the extent of the colour gamut and the perceptual isometry (visual uniformity) of the colour space. For a particular bit-depth, quantization error is generally inversely related to the size of the colour gamut since the same number of code values will be distributed within a larger colour gamut volume. Hence, the importance of the colour gamut being no larger than necessary for a particular application. Colour spaces that are less visually uniform will have wider distributions of colour differences and larger maximum colour differences.

The threshold for acceptable quantization errors will be application dependent. In some applications, the threshold may be any colour difference that is visually detectable, while other applications may be less demanding. In cases where aggressive image-processing algorithms will be applied (e.g. tone scale manipulations, chroma-boosting transforms or image compression), it will be more critical that quantization errors be minimized since these errors can be amplified during processing. Applications having low-noise image sources (e.g. computer-generated gradients) will have a higher level of sensitivity to quantization errors.

C.7 Quantization efficiency

Another useful criterion for integer digital encodings is to calculate the quantization efficiency, which is the colour encoding bit-depth needed to achieve an acceptable level of image quality. One straightforward means for determining the quantization efficiency is to compute quantization error statistics at a series of different bit-depths. The bit-depth corresponding to a threshold level of the quantization error statistic can then be determined. As discussed in C.6, the threshold for acceptable quantization errors will be application dependent. In cases where it is desired to limit the worst-case error, a threshold can be set for the maximum colour difference. In cases where it is desired to meet an average image quality requirement, a threshold can be set for the average colour difference.

C.8 Perceptual isometry

The perceptual isometry (visual uniformity) of a colour encoding is another criterion that may be important in some applications. A perceptually isometric colour encoding is one where differences of a fixed magnitude in the colour encoding (equal distances in its colour space) are equally perceptible throughout colour space. Perceptually isometric colour encodings may be preferred for any of several reasons. Algorithms for making aesthetic or editorial changes to images are also often designed to operate optimally in a perceptually isometric colour space so that an image manipulation will have approximately equal impact throughout the colour space and be less likely to introduce discontinuities or other artefacts. Also, as discussed in C.6 and C.7, perceptual isometry is closely related to the visual uniformity of quantization errors throughout the colour space, which in turn affects the quantization efficiency. Metrics can be defined to measure the overall perceptual isometry of a colour encoding. Metrics can also be developed to characterize the visual uniformity with respect to individual colour dimensions such as lightness, chroma and hue.

It should be noted that considerable evidence exists that those colour spaces commonly regarded as being perceptually isometric are only approximately so. (See Reference [12].) Therefore, the development of metrics to characterize perceptual isometry can be somewhat problematic since no perfect colour appearance model has yet been developed that can be used as a basis for comparison. Any metric based on evaluation using an existing colour appearance model will tend to build any inherent distortions associated with that particular colour appearance model into the metric. While these distortions may be small for the case of characterizing lightness uniformity, they can be significant when dealing with more complex attributes such as hue uniformity. Also, when appearance models are properly used, perceptual isometry will be viewing condition dependent.

C.9 Complexity of conversions to other important colour encodings

There will be a need in many digital photography and graphic technology applications to convert images in a particular colour encoding to some other colour encoding(s) as the image progresses through an imaging chain and is used for various purposes. One important example is the formation of a preview image that can be displayed on a CRT. The complexity of the transformations needed to convert the image to particular colour encodings that are relevant for a particular application is an important colour encoding selection criterion, particularly for those workflows where real-time image display is critical. Some colour encodings that will be important in various applications include sRGB, ICC PCS XYZ, ICC PCS LAB and SWOP CMYK. The computation time that it will take to convert from a given colour encoding to one of these important colour encodings (or to some other important colour encoding) will depend on many factors including the processing hardware, the software design, etc. As a result, it is difficult to define a meaningful quantitative metric for this criterion. However, it is possible to classify the complexity of these colour transformations into a number of coarse qualitative categories such as:

— 1-D LUT;

— (1) 3×3 matrix and (1 or 2) 1-D LUTs;

— (2) 3×3 matrices and (1, 2 or 3) 1-D LUTs;

— 3-D LUT or other complex transform.

It may also be important to consider related questions such as whether lossless invertability can be maintained when the colour transformation is applied to an image. This can be affected by factors such as mismatches in colour space white points, colour gamuts or colour component transfer functions.

It should be noted that for scene-referred colour encodings, the transformation to an output-referred colour space, such as sRGB, will necessarily include a colour-rendering operation. Therefore, the complexity of this transformation will depend on the particular form of the colour-rendering transform. For purposes of evaluating the transform complexity, a colour-rendering function typical of that which will be used in the particular application should be used when it is known. For cases where the typical rendering transform is not known, the simplest colour-rendering function should be assumed. For RGB colour encodings, this generally would correspond to a tone scale transform implemented as a 1-D LUT.

C.10 Compressibility

Image compression is an important step in many digital photography and graphic technology workflows in order to minimize memory requirements for image storage, as well as bandwidth requirements for image transmission. Typical “lossy” compression techniques, such as JPEG and JPEG2000, exploit the fact that the human visual system has a reduced sensitivity to chrominance errors as compared luminance errors, and high frequency errors as compared to low-frequency errors. Different colour encodings may be more or less conducive to being compressed with a particular compression algorithm. For example, many compression algorithms will tend to distribute errors relatively uniformly in code value space (i.e. the standard deviation of the errors in the shadow regions will be the same number of code values as that in the highlight regions). This suggests a general principle that colour encodings with a higher degree of perceptual isometry will be better adapted for image compression. For example, colour spaces that are linear with respect to radiance (e.g. CIE XYZ) will tend to produce errors that are more visible in darker image regions than lighter image regions. This is because a given code value error will produce a larger visual change for a darker colour than a lighter colour. (This is the main reason that many colour encodings include a nonlinear digital encoding scheme to more closely approximate the human visual response.)

There are several metrics that could be calculated to assess the compressibility for a colour encoding. One useful metric is to determine the image quality level achieved for a constant compression ratio for some set of typical test images. Another useful metric is to determine the average compression ratio corresponding to some threshold level of image quality. The image quality of the compressed/decompressed images is a key component for both of these metrics. Unfortunately, this is rather difficult to determine quantitatively. Simple metrics, such as the r.m.s. error between the original and compressed/decompressed images, do not correlate well with perceived image quality since they fail to account for the characteristics of the human visual system. Various visual difference metrics have been proposed in recent years based on visual system models of varying complexity. However, a widely accepted metric has not been identified that can be used to predict the quality loss associated with image artefacts such as compression errors. (s-CIELAB is one example of a relatively simple visual system model that could be used as the basis of a visual difference metric.) At present, the most accurate image compression quality assessments require the use of subjective psychophysical experiments using many images and judges. Such experiments can be time-consuming and costly to carry out.

The relative compressibility of various colour encodings will be a function of the compression method, as well as the image quality requirements and the characteristics of the image population for a particular application.

It should be noted that most image compression algorithms for RGB images include a step of rotating the RGB image to a luma-chroma space (e.g. YC_rC_b) before compressing the image. Thus, when evaluating compressibility criteria, it will usually be appropriate to apply such rotations for colour encodings based on RGB colour space metrics. Some compression algorithms also support the application of a nonlinear colour component transfer function before compressing the image. In such cases, it may also be appropriate to incorporate a nonlinear encoding to improve the compressibility of linear colour encodings.

C.11 Compatibility with existing industry practice

When evaluating colour encodings, it is also important to consider how they fit within the context of expected workflows and applications. One important criterion is compatibility with software/hardware platforms that may be fundamental to a particular market. For example, certain software applications may be incompatible with colour encodings requiring more than 8-bits/channel, or colour encodings that use signed integer representations.

Another important criterion is compatibility with important image-processing operations/algorithms. In many cases, these algorithms have been designed with certain fundamental assumptions about the characteristics of the colour encoding. For example, certain tone scale manipulation operations may be incompatible with colour encodings utilizing signed colour values. Similarly, transparency operations may be incompatible with over-ranged colour values, and many computer-generated graphics algorithms may require a colour encoding that is linear with image radiance. These requirements will generally be application specific.

C.12 Hue shifts induced by nonlinear tone scale manipulations

Many common image-processing operations are based on the application of nonlinear transforms to one or more of the individual colour channels of an RGB image. For example, brightness, contrast and colour balance adjustments are all commonly implemented in this fashion. These nonlinear transforms will, in general, modify the relative ratios of the red, green and blue channel data. This can lead to unwanted hue shifts, particularly for high chroma colours. Hue shifts are particularly problematic in reproductions of natural chroma gradients, having constant hue and saturation. Such gradients tend to occur when rounded surfaces are illuminated by a moderately directional light source. In such situations, chroma increases with distance from the specular highlight and then decreases again as the shadows deepen. The definition of the colour encoding can have a significant effect on the hue shift characteristics. The objective when evaluating colour encodings is to eliminate or minimize objectionable hue shifts at the expense of less noticeable or less likely hue shifts.

Hue shifts introduced by the application of nonlinear transformations can be studied by applying an aggressive nonlinear tone scale to chroma series generated for a set of colour patches. It is generally desirable to include primary colours such as red, green, blue, cyan, magenta and yellow, as well as other important colours such as skin tones. The hue differences between the input and output colour values can then be computed using a uniform colour space or an appropriate colour appearance model. Hue shifts in skin tones and yellows, particularly in the direction of green, are considered to be the most objectionable in many cases.

Bibliography

[1] ANSI/I3A IT10.7466, *Photography — Electronic still picture imaging — Reference input medium metric RGB color encoding (RIMM-RGB)*

[2] ANSI/I3A IT10.7666, *Photography — Electronic still picture imaging — Reference output medium metric RGB color encoding (ROMM-RGB)*

[3] ICC.1:2001-12 *File format for color profiles (Version 4.0.0)*

[4] IEC 61966-2-1:1999, *Multimedia systems and equipment — Colour measurement and management — Part 2-1: Colour management — Default RGB colour space — sRGB*

[5] IEC 61966-2-1 A1:2003, Amendment 1, *Multimedia systems and equipment — Colour measurement and management — Part 2-1: Colour management — Default RGB colour space — sRGB*

[6] IEC 61966-2-2:2003, *Multimedia systems and equipment — Colour measurement and management — Part 2-2: Colour management — Extended RGB colour space — scRGB*

[7] ISO 3664:2000, *Viewing conditions — Graphic technology and photography*

[8] ISO 17321-2:—[1)], *Graphic technology and photography — Colour characterisation of digital still cameras (DSCs) — Part 2: Methods for determining transforms from raw DSC to scene-referred image data*

[9] ITU-R BT.601-5, *Studio encoding parameters of digital television for standard 4:3 and wide-screen 16:9 aspect ratios*

[10] ITU-R BT.709-3, *Parameter values for the HDTV standards for production and international programme exchange*

[11] I3A 7667, *Electronic Still Picture Imaging — Extended sRGB color space — e-sRGB*

[12] Braun, G.J. Fairchild, M.D. and Ebner, F. Color gamut mapping in a hue-linearized CIELAB color space. in *Proceedings of Sixth Color Imaging Conference*, 1998, pp. 163-158

[13] Giorgianni, E.J. and Madden, T.E. *Digital color management: encoding solutions*, Addison-Wesley, Reading, MA, 1998

[14] Holm, J. Issues relating to the transformation of sensor data into standard color spaces, in *Proceedings of Fifth Color Imaging Conference*, 1997, pp. 290-295

[15] Holm, J. Edwards, E. and Parulski, K. Extended colour encoding requirements for photographic applications. in *Proceedings of the CIE Expert Symposium 2000 on Extended Range Colour Spaces,* CIE x021:2001, pp. 16-25

[16] Hubel, P.M. The perception of color at dawn and dusk, in *Proceedings of Seventh Color Imaging Conference*, 1999, pp. 48-51

[17] Jones, A. and Condit, H.R. Sunlight and skylight as determinants of photographic exposure, *J. Opt. Soc. Am.*, **38**, 1948, pp. 123-178 and **39**, 1949, p. 94

[18] Pointer, M.R. The gamut of real surface colours. *Color Res. Appl.,* **5**, 1980, pp. 145-155

[19] Spaulding, K.E. Image states and standard extended gamut color encodings. in *Proceedings of the CIE Expert Symposium 2000 on Extended Range Colour Spaces*, CIE x021:2001, pp. 59-70

1) Under preparation.

[20] SPAULDING, K.E., WOOLFE, G.J. and GIORGIANNI, E. J. Optimized extended gamut color encodings for scene-referred and output-referred image states. *J. Imaging Sci. Technol*., **45**, 2001, pp. 418-426

[21] SUSSTRUNK, S., BUCKLEY, R. and SWEN, S. Standard RGB color spaces, in *Proceedings of Seventh Color Imaging Conference*, 1999, pp. 127-134

[22] ZHANG, X. and WANDELL, B.A. A spatial extension of CIELAB for digital color-image reproduction. *J. SID*, **5**, 1997, pp. 61-63

[23] ISO 13655, *Graphic technology — Spectral measurement and colorimetric computation for graphic arts images*

ICS 37.040.99; 37.100.01

Price based on 47 pages