
**Graphic technology and
photography — Colour
characterization of digital still
cameras (DSCs) —**

**Part 2:
Considerations for determining scene
analysis transforms**

*Technologie graphique et photographie — Caractérisation de la
couleur des appareils photonumériques —*

*Partie 2: Considérations pour déterminer les transformations
d'analyse de scène*



STANDARDSISO.COM : Click to view the full PDF of ISO/TR 17321-2:2012



COPYRIGHT PROTECTED DOCUMENT

© ISO 2012

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

| | Page |
|---|-----------|
| Foreword..... | iv |
| Introduction..... | v |
| 1 Scope..... | 1 |
| 2 Definitions..... | 1 |
| 3 Goals..... | 4 |
| 4 Fundamental colour-related DSC characteristics..... | 5 |
| 4.1 Camera gain..... | 5 |
| 4.2 Camera dark current..... | 5 |
| 4.3 Focal plane opto-electronic conversion function (FP OECF)..... | 6 |
| 4.4 Camera opto-electronic conversion function (Camera OECF)..... | 6 |
| 4.5 Camera flare average percent..... | 6 |
| 4.6 Camera spectral sensitivities (including non-removable optical elements)..... | 6 |
| 4.7 Removable optical element spectral transmittances..... | 6 |
| 5 Scene analysis transform parameters..... | 7 |
| 5.1 Determined using spectral measurements..... | 7 |
| 5.2 Determined using test targets..... | 8 |
| 6 Scene analysis transform determination using spectral measurements..... | 9 |
| 7 Scene analysis transform determination using test targets..... | 10 |
| 7.1 Procedure..... | 10 |
| 7.2 Test target characteristics..... | 12 |
| 7.3 Capturing images of test targets..... | 12 |
| 8 Applying the scene analysis transform and encoding..... | 13 |
| 9 Considerations for selecting scene analysis transforms..... | 14 |
| 10 Suggested metadata for scene-referred colour encodings..... | 14 |
| 10.1 Camera characteristics..... | 14 |
| 10.2 Colour encoding characteristics..... | 14 |
| 10.3 Image specific characteristics..... | 14 |
| Annex A (informative) Example calculation results..... | 15 |
| Bibliography..... | 17 |

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.

In exceptional circumstances, when a technical committee has collected data of a different kind from that which is normally published as an International Standard ("state of the art", for example), it may decide by a simple majority vote of its participating members to publish a Technical Report. A Technical Report is entirely informative in nature and does not have to be reviewed until the data it provides are considered to be no longer valid or useful.

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

ISO/TR 17321-2 was prepared by Technical Committee ISO/TC 42, *Photography*.

ISO 17321 consists of the following parts, under the general title *Graphic technology and photography — Colour characterization of digital still cameras (DSCs)*:

- *Part 1: Stimuli, metrology and test procedures*
- *Part 2: Considerations for determining scene analysis transforms*

Introduction

Digital still cameras (DSCs) have become the predominant means of photographic image capture, but the nature of the image data produced by different cameras, or even by the same camera operating in different modes, is quite variable. This variability can cause problems in workflows, miscommunications and interoperability issues.

This Technical Report provides information about methods for determining scene analysis transforms, which are transforms that convert raw image data to scene-referred image data. This information is provided in the form of a Technical Report because there are a number of choices to be made when determining scene analysis transforms. These choices are influenced by the subject matter being photographed (including the scene illumination), the scene adopted white and the adopted white of the scene-referred colour encoding to be used, aesthetic choices regarding scene analysis colour error minimization, and other considerations. It is not possible to provide more specific recommendations because the spectral responses of DSC colour analysis channels do not, in general, match those of a typical human observer, such as defined by a CIE standard colourimetric observer. Nor do the responses of different DSCs ordinarily match each other. This Technical Report outlines considerations relevant to the determination of scene analysis transforms based on the minimization of errors in specified colour spaces. The DSC characterization data obtained using ISO 17321-1 serve as the raw DSC image data.

Good understanding of this Technical Report requires that the three fundamental modes of DSC operation be distinguished: the raw mode, the scene-referred mode, and the output-referred mode. When operating in the raw mode, a DSC records image data that is most closely related to the sensor response. Some types of processing may have been performed, such as dark current subtraction, defect removal and colour filter array interpolation, but neither a scene analysis nor a colour rendering transform has been applied. Any encoding transform typically consists of only a non-linearity to better align the quantization intervals with the image noise characteristics, and possibly some form of compression.

When operating in the scene-referred mode, a DSC records image data that represents an estimate of the scene or focal plane image relative colourimetry, typically with white balancing to the encoding adopted white. The image data has not undergone colour rendering for some anticipated output medium and viewing conditions. In order to produce output-referred images intended for reproduction, it is necessary to either apply a colour rendering transform directly to the scene-referred images, or convert them to a working colour space where the desired colour rendering is applied. Camera controls or raw processing software can offer some aesthetic choices when converting to scene-referred, but the results of such choices need to be viewed through the intended colour rendering transform in order to see their effect on the final output. The image data are encoded prior to applying the colour rendering transform, and are therefore not an encoding of the intended output colourimetry. At present, few DSCs offer an in-camera scene-referred mode, although some camera raw processing applications have this capability.

When operating in the output-referred mode, the DSC controls are set to achieve the desired output directly, thereby incorporating the colour rendering, and in many cases a reference output device encoding (such as for a CRT monitor) in the image file. When operating in this mode the DSC encodes the colourimetry of the intended output on the reference medium, not scene-referred colourimetry. Also, the output-referred colourimetry can be in different encodings, with different reference media, and in some cases will need to be colour re-rendered and/or re-encoded to produce different reproductions.

The information provided in this Technical Report is intended to help camera and raw processing software manufacturers, professional photographers and colour measurement applications to determine, communicate about, and select DSC scene analysis transforms. However, it will often not be practical for end users to determine scene analysis transforms themselves. In addition to the requirement for raw DSC image data, relatively sophisticated and expensive measurement equipment is required to obtain chart patch spectral reflectance or radiance, illumination source spectral power, and DSC spectral sensitivity (as described in ISO 17321-1).




Figure 1 — Generic image workflow for digital photography

STANDARDSISO.COM : Click to view the full PDF of ISO/TR 17321-2:2012

Graphic technology and photography — Colour characterization of digital still cameras (DSCs) —

Part 2: Considerations for determining scene analysis transforms

1 Scope

This Technical Report provides information about methods for determining scene analysis transforms based on the minimization of errors in estimated scene or focal-plane colourimetry, including corresponding colourimetry. These transforms are limited in applicability to raw DSC image data.

This Technical Report concerns only the creation and encoding of scene-referred and focal-plane-referred image data. It does not address the encoding of output-referred image data. It also does not provide information relating to the specification of metadata items describing intended artistic adjustments, colour rendering and viewing.

This Technical Report does not address how to choose adopted white points or how to process scene-referred image data to produce output-referred image data.

This Technical Report is not intended to be comprehensive or complete; it is an overview intended to enable improved practices and communications.

2 Definitions

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

2.1

absolute colorimetric coordinates

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

EXAMPLE If the colorimetric coordinates used are CIE 1931 standard 2° observer tristimulus values, the Y value should correspond to the luminance, not the luminance factor (or some scaled value thereof).

2.2

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 1 The adapted white can vary within a scene.

NOTE 2 No assumptions should be made concerning the relation between the adapted 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 can affect perception. It is easy to arrange conditions for which a near perfectly reflecting diffuser will appear to be grey or coloured.

NOTE 3 See adapted white (2.3).

2.3

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 can vary within a scene, if such variation is supported by the imaging system.

NOTE 2 The adopted white is not required to be an estimate or approximation of the adapted white. For example, if a scene lit by tungsten illumination is captured using a DSC with the white balance set to D55 (daylight), the adopted white will be D55 but the adapted white will be closer to a tungsten illuminant (e.g. ISO 7589 Studio Tungsten or CIE Illuminant A).

NOTE 3 See adapted white (2.2).

2.4

scene analysis transform spectral limit

mapping of spectral (monochromatic) colours as captured by a DSC and transformed to a scene-referred colour encoding using a specified scene analysis transform

2.5

colour component transfer function

CCTF

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 non-linear 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 non-linear functions such as a power-law (i.e. "gamma") function or a logarithmic function. However, in some cases a linear colour component transfer function will be used.

2.6

colour gamut

solid in a colour space, consisting of all those colours that are either: 1) present in a specific scene, artwork, photograph, photomechanical or other reproduction; 2) capable of being created using a particular output device and/or medium

NOTE See scene analysis transform spectral limit (2.4).

2.7

colour matching functions

tristimulus values of monochromatic stimuli of equal radiant power

[CIE Publication 17.4, 845-03-23]

2.8

colour pixel reconstruction

algorithm that creates a fully populated colour image record from the output of a colour filter array type sensor by interpolating values for each colour at each pixel location, also known as demosaicing or colour pixel interpolation

2.9

colour rendering

mapping of image data representing the colorimetric coordinates of the elements of a scene or original to image data representing the colorimetric 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.

2.10**colour space**

geometric representation of colours in space, usually of three dimensions

[CIE Publication 17.4, 845-03-25]

2.11**corresponding colorimetry**

colorimetric coordinates for samples that produce a visual match with different viewer adaptation states

EXAMPLE Chromatic adaptation transforms estimate corresponding colourimetry for viewing conditions where only the adapted white chromaticity is different. Colour appearance models can be used to estimate corresponding colourimetry where more aspects of the viewing conditions are different. The colour appearance model is used to calculate appearance correlates for one viewing condition. The inverse of the colour appearance model is then used to calculate colourimetric coordinates for a second viewing condition from the colour appearance correlates.

2.12**digital still camera****DSC**

device which incorporates an image sensor and which produces a digital signal representing a still picture

NOTE A digital still camera is typically a portable, hand-held device. The digital signal is usually recorded on a removable memory, such as a solid-state memory card or magnetic disk.

2.13**output-referred image data**

image data which represents the colorimetric 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 The 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 2 Output-referred image data can 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.

2.14**raw DSC image data**

image data produced by or internal to a DSC that has not been processed, except for A/D conversion and the following optional steps: linearization, dark current/frame subtraction, shading & sensitivity (flat field) correction, flare removal, white balancing (e.g. so the adopted white produces equal RGB values or no chrominance), missing colour pixel reconstruction (without colour transformations)

NOTE See scene-referred image data (2.17).

2.15**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 can represent an actual view of the natural world or a simulation of such a view.

2.16**scene analysis transform**

transform that converts raw DSC image data to scene-referred image data

2.17

scene-referred image data

image data which represents estimates of the colorimetric coordinates of the elements of a scene

NOTE 1 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, except possibly in a special mode intended for this purpose. Typically, DSCs write standard output-referred image data where colour rendering has already been performed.

NOTE 2 Scene-referred image data typically represents relative scene colourimetry estimates. Absolute scene colourimetry estimates can 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 3 Scene-referred image data can 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 4 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 can be correctable if the transformation used to produce the scene-referred image data are known, and the colour encoding used for the incorrect scene-referred image data has adequate precision and dynamic range.

NOTE 5 The scene can correspond to an actual view of the natural world, or a simulation of such a view. It can 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.

2.18

tristimulus values

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]

NOTE See colour matching functions (2.7).

2.19

working colour space

colour space encoding in which operations such as image edits, enhancements or colour rendering are performed

NOTE 1 The image state in a working colour space can change as operations are performed.

NOTE 2 If operations performed in a working colour space are guided by viewing the image on a medium, that medium and the associated viewing conditions become the reference for the resulting image.

NOTE 3 See colour matching functions (2.7).

3 Goals

The goals of this Technical Report are as follows.

- To list the fundamental colour-related characteristics of DSCs.
- To document some methods and parameters used for the conversion of raw DSC image data to scene or focal plane colourimetry estimates
 - for the case where scene analysis transforms are determined using spectral measurements, and

- for the case where scene analysis transforms are determined by capturing test targets.
- To provide recommendations for specifying the encoding of scene and focal plane colourimetry estimates.
- To provide recommendations for metadata to be included in scene-referred image files or other specified locations, communicating the colour-related characteristics of the capture device, the scene analysis transform used, and the colour encoding used.

4 Fundamental colour-related DSC characteristics

4.1 Camera gain

Camera gain is the ratio of the digital count obtained to the sensor exposure for the linear region of the sensor and analogue-to-digital converter response range.

- The sensor exposure can be either a photometric exposure or an integrated channel radiometric exposure, but different gains will typically result for each type of exposure.
- Neutral balance is achieved by adjusting the channel gains, either in analogue or digital processing.
- If the sensor or the encoding of the sensor signals is non-linear, the gain will be a function of the sensor exposure.
- Even for a sensor with inherently linear response, a non-linear encoding can be used to take advantage of noise statistics to reduce the bit-depth required to encode the sensor signals.
- In use, digital cameras can offer selectable gains or apply automatic gain control in attempt to maintain normal signal levels when different Exposure Indexes (EIs) are used. Likewise, selectable or automatic channel gains are typically used to achieve a satisfactory neutral balance. For the characterization of raw mode camera behaviour, the overall and channel-specific camera gains need to be fixed, i.e. any automatic gain and white balance needs to be disabled.
- Generally, automatic gain control is not applied when a camera is set to raw mode operation, to provide the maximum capture dynamic range. It is assumed that any desired overall gain or neutral balance will be applied later.
- Different gains applied to address the use of different EIs or neutral balances are applied to linear sensor signals, prior to any non-linear encoding.

4.2 Camera dark current

Camera dark current is the digital counts recorded in the absence of any sensor exposure.

- Dark current is typically a function of the sensor temperature and integration time.
- The dark current value can be an average over the sensor (or some area of the sensor), or can be determined on a per-pixel basis using dark frame exposures. However, it should be noted that since dark current is noisy, it is highly desirable to average multiple dark frames.
- Many cameras perform dark current subtraction in the analogue domain, before digitization, in which case the average digital dark current can be close to zero.
- Use of an encoding that supports negative values preserves the statistics of the dark current after subtraction.

4.3 Focal plane opto-electronic conversion function (FP OECF)

The FP OECF is the functional relationship between the digital count obtained and the sensor photometric exposure for the selected adopted white spectral radiance distribution.

- In order to account for the spectral characteristics of the optics, the FP OECF should be determined with the camera lens, filters and other optical elements in place, i.e. using the alternative focal plane OECF determination method specified in ISO 14524.
- If the dark current is not subtracted before encoding, it will be included in the FP OECF.

4.4 Camera opto-electronic conversion function (Camera OECF)

The Camera OECF is the functional relationship between the digital count obtained and the scene luminance (or luminance factor) for the selected adopted white spectral radiance distribution.

- The difference between the Camera OECF and the FP OECF usually results from lens flare.

4.5 Camera flare average percent

For a particular camera configuration, the camera flare average percent is the typical average flare exposure as a percent of the mean exposure for each channel.

- In some workflows camera flare information is not required; manual compensation is performed as needed or a typical value for the type of camera being used is assumed.

4.6 Camera spectral sensitivities (including non-removable optical elements)

The camera spectral sensitivities are the functional relationship between the wavelength of the exposing radiation and the resulting linearized digital count for each channel, including compensation for the power of exposing radiation at each wavelength.

- When capturing narrow-band illumination with the camera to determine the spectral sensitivities, the exposure settings need to be chosen to avoid clipping on any channel (while providing reasonable signal levels) and then should not be changed (to minimize the effects of aperture and shutter setting variations).
- The camera spectral sensitivities need to be determined after dark current subtraction, and in the case of a non-linear response after FP OECF inversion.
- The exposing radiation power compensation can be achieved by either adjusting the radiation power, or in the case of linear (or linearized) camera response by adjusting the digital values consistent with the measured incident power.

4.7 Removable optical element spectral transmittances

The spectral transmittances of any removable optical elements, such as lenses and filters, used on the camera, or in the case of a standard filter with published spectral transmittances, the filter designation.

- In many cases, acceptable results can be obtained if the ISO 7589 standard lens transmittances are assumed.

5 Scene analysis transform parameters

5.1 Determined using spectral measurements

When scene analysis transforms are determined using spectral measurements, the following parameters are used.

- Tristimulus value definition (e.g. CIE 1931 2° observer)
- Adopted white spectral radiance distribution

NOTE The adopted white spectral radiance distribution need not be the same as that of the scene illumination, or any object in the scene. The adopted white spectral radiance distribution is the spectral power distribution that is considered to be neutral in the scene and converted to the adopted white of the scene-referred colour encoding.

- Training spectra spectral radiances and tristimulus values under the scene illumination
 - It is generally preferable to choose training spectra that simulate the spectral radiances of real-world colours of interest. If these are not available monochromatic stimuli can be used to ensure the entire spectrum is reasonably well analysed.

NOTE Standard databases of object spectra often do not represent real-world spectra well because the measurements are not taken *in situ*. The actual spectral radiances of objects in scenes are strongly influenced by locally transmitted, scattered, and reflected light, and can change substantially if the objects are removed from their environment. For example, a blade of grass in a lawn is much greener than an isolated one because the local illumination in the lawn is very green due to light from the other blades of grass.

- Scene-referred colour encoding adopted white chromaticity (if different from scene adopted white chromaticity) and white balance transform applied to training spectra tristimulus values to produce aim tristimulus values
- Form of the scene analysis transform (e.g. 3x3 matrix) to the scene analysis colour space
 - The scene analysis colour space needs to be linear with respect to scene radiance or focal plane irradiance, e.g. CIE XYZ with some specified adopted white or some linear colourimetric RGB.
- Error minimization
 - Colour space, including the transform from the scene analysis colour space to the error minimization colour space.
 - If monochromatic colours are used, it is recommended to use an RGB colour space with a non-linear colour component transfer function for better perceptual uniformity, such as a square root function.
 - Research has shown^[21] that monochromatic primaries with wavelengths of 450, 540 and 620 NM work well, although other primaries such as ITU-R BT.709 or RIMM RGB primaries can be used. Different primary choices will usually produce different results making the selection of the error minimization colour space an aesthetic choice.
 - If a colour space is used where the colour space values are determined relative to a specified white point (such as with CIE L*a*b* or CIE L*u*v*), both the aim and camera colour space values must be normalized to be relative to the same colour space white point. Typically this white point will be substantially lower than the camera white clip in order to provide headroom for the capture of above-white colours.
 - error minimized (e.g. least square error)
 - training spectra weights (if not all equal)

- constraints (generally it is a good idea to enforce neutral preservation when the scene analysis transform is determined).

NOTE If the camera spectral sensitivities are colour matching functions, different training spectral radiances and error minimization parameters will produce the same scene analysis transform.

5.2 Determined using test targets

When scene analysis transforms are determined using test targets, the following parameters are used.

- Tristimulus value definition (e.g. CIE 1931 2° observer)
- Scene adopted white chromaticity
- Target patch spectral radiances and tristimulus values under scene illumination (spectral radiances are only used to verify appropriateness of patches)
 - It is generally preferable to choose test patches that simulate the spectral radiances of real-world colours of interest. Note that there is no benefit to choosing test patches that only simulate the colourimetry of real-world objects of interest (unless the camera spectral sensitivities are colour matching functions).
 - Test patches that are created using a limited set of colorants (such as those produced using conventional photographic or printing processes) should be avoided when creating scene analysis transforms to be applied to images of real-world scenes. If such patches are used, the scene analysis transforms determined will be highly tuned to the spectral characteristics of the colorants used and can produce unacceptable analysis errors for objects with substantially different spectral characteristics.
- Destination adopted white chromaticity (if different from scene adopted white chromaticity) and white balancing transform applied to patch tristimulus values to produce aim tristimulus values
- Scene analysis transform form (e.g. 3x4 matrix) to the scene analysis colour space
 - The scene analysis colour space is typically linear with respect to scene radiance or focal plane irradiance, e.g. CIE XYZ with some adopted white or some linear colourimetric RGB.
- Camera flare handling
- Error minimization colour space (e.g. CIE $L^*u^*v^*$ or CIE $L^*a^*b^*$)
 - Including the transform from the scene analysis colour space to the error minimization colour space.
 - If a colour space is used where the colour space values are determined relative to a specified white point (such as with CIE $L^*u^*v^*$ or $L^*a^*b^*$), both the aim and camera colour space values must be normalized to be relative to the same colour space white point. Typically this white point will be substantially lower than the camera white clip in order to provide headroom for the capture of above-white colours.
- Error metric (e.g. mean or maximum ΔE^*uv or CIE ΔE^*ab)
- Patch weights (if not all equal)
- Constraints (e.g. neutral preserving)

NOTE If the camera spectral sensitivities are colour matching functions, different target patch spectral radiances and error minimization parameters will produce the same scene analysis transform.

6 Scene analysis transform determination using spectral measurements

Scene analysis transforms can be determined using spectral measurements by performing the following steps:

- a) Adapt the training spectra tristimulus values to be relative to the destination adopted white by applying the white balancing transform (if necessary).
- b) Convert the white balanced aim tristimulus values to the error minimization colour space.
- c) Calculate the camera colour channel responses to the training spectra by first multiplying the camera channel spectral sensitivity, the removable optical element spectral transmittance, and the training spectral radiance vectors, and then summing the resulting vectors to obtain a single linear response value for each channel for each training spectrum.
- d) Apply the channel gains necessary to produce the selected adopted white to the calculated linear response values to produce white balanced linear response values.
- e) Using regression (or iteration, if necessary), determine the scene analysis transform that, when applied to the neutral-balanced linear response values produces the smallest weighted error with respect to the aim values after the neutral-balanced linear response values are converted to the error minimization colour space, using the selected error metric and subject to the selected constraints.

In Figure 2 below, the transform for the image on the top left used CIE XYZ as the error minimization colour space, with least-squares error minimization of spectral training spectra and equal weighting. The transform for the image on the top right used a non-linear RGB error minimization colour space, with least squares error minimization of monochromatic training spectra and illuminant power weighting. The transform for the image on the lower left used CIE L*u*v* as the error minimization colour space with ΔE^*_{uv} error minimization of the Macbeth ColourChecker™ patches, weighted equally. The image on the lower right is a virtual colour checker with the aim patch values for comparison. All images were converted to sRGB colourimetrically for inclusion in this Technical Report. Note that the captured images contain some uncorrected illumination falloff.



Figure 2 — Scene-referred images of a Macbeth ColourChecker™¹⁾ produced from the same raw image data but using different scene analysis transforms

7 Scene analysis transform determination using test targets

7.1 Procedure

Scene analysis transforms can be determined using test targets by performing the following steps:

- a) Measure the tristimulus values of the test target under the scene illumination from the camera position using a telescopic instrument.
 - The acceptance angle of the instrument should be small enough, and the test target patches large enough so each patch can be measured individually with some margin for alignment error.
 - Measurement instruments with an angle of acceptance between 0.5° and 2° are typically used.
 - The edges of the patches should be avoided when taking measurements.
 - Care should be taken to minimize and/or correct for measuring instrument flare. One way to achieve this is to use a black mask that covers the chart and surrounding area except for the patch being measured. However, care should then be taken to ensure the mask does not obstruct any illumination falling on the patch being measured.

1) Macbeth ColourChecker™ is an example of a suitable product available commercially. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO of this product.

- In non-critical applications, it can be acceptable to measure the test target patches using a different measurement geometry, but this will often result in incorrect measurements because of viewing flare. Likewise, it can be acceptable to measure the test target patch spectral reflectances or transmittances and calculate tristimulus values for the illumination, but this can produce incorrect values when the test target has patches that exhibit fluorescence.

NOTE Errors in scene analysis, especially those that result in saturation increases, can sometimes appear more pleasing when viewed directly, but should generally be avoided because colour rendering transforms will assume a preferential saturation boost has not been incorporated into the scene-referred image data.

- b) Adapt the test patch tristimulus values to be relative to the destination adopted white by applying the white balancing transform (if necessary).
- c) Convert the white balanced aim tristimulus values to the error minimization colour space.
- d) Capture an image or images of the test target under the scene illumination.
 - The camera data should be “raw” or native camera RGB. Cameras that deliver Rec 709 data are performing internal colour processing to a CRT display, producing output-referred images, the use of which is not recommended.
 - Image data can be linear, or encoded using a function such as a log or gamma. Ideally, linear image data should be 12-bits per channel or more. Log or gamma image data should be 8-bits per channel or more.
 - To avoid colour shifts, data compression, if used, should be near lossless. The quantization used in some lossy compression methods can cause minor colour shifts.
 - The image data needs be saved prior to any colour conversions or adjustments.
 - The following should be recorded or included in the image metadata:
 - The test target make, model and serial number
 - The camera make, model and serial number
 - The camera settings needed to be able to set the camera to exactly the same operation
 - Any lenses and filters used
 - The lights, gels and power levels used
- e) Average multiple pixels in each patch, and potentially from the same patch in multiple images, to reduce noise and average out other sources of variability. Pixels near the edges of the patches should be excluded.
- f) Apply the inverse of the Camera OECF to the average patch values to produce neutral-balanced scene-referred linear response values.
 - The Camera OECF should be determined according to ISO 14524, using a test target with a dynamic range that spans that of the colour characterization test target. If the camera lens falloff is corrected, it can be acceptable to use other test targets for camera OECF determination, including neutral patches incorporated in the colour characterization test target.
 - The illumination and neutral balance used when determining the camera OECF should also be used for the camera characterization, but slight neutral balance differences can be addressed by also regressing separate channel gains for neutral patches on the colour characterization test target (assuming both the OECF and colour characterization neutral patches are truly neutral).
- g) Using regression (or iteration, if necessary), determine the scene analysis transform that, when applied to the scene-referred linear response values produces the smallest weighted error with

respect to the aim values after the linear response values are converted to the error minimization colour space, using the selected error metric and subject to the selected constraints.

- Generally it is recommended to include an offset in the matrix to address any differences in dark current and flare from the camera when the OECF was determined.
- With test target characterization, maximum error can be a better choice for minimization than mean error because mean error minimization can result in unacceptably large errors for some colours, while maximum error minimization will typically have less effect on the mean error.

7.2 Test target characteristics

- Test targets can be reflective or transmissive. For pure camera characterization, a transmissive test target is recommended. A transmissive test target can produce a much higher dynamic range for determining the OECF, and a larger colour gamut. For in-the-field calibration and characterization, a reflective test target is recommended. The reflective test target can more easily capture the actual scene illumination, but will usually have a smaller dynamic range and colour gamut.
- Test targets should use colorants (dyes or pigments) for each patch, with each colorant having its own unique spectral characteristics (reflectance or transmittance). Ideally these spectral characteristics should be similar to those of important colours to be captured. When patches are made from a small number of colorants, the effective number of colour samples is only the small number of colorants.
- With transmissive targets, scattering should be minimized which can favour the use of dyes over pigments.
- Test targets need to be optically stable; resistant to fading or other changes caused by illumination, ultraviolet or infrared, temperature and humidity.
- Test targets should not affect the polarization of the illumination or emit polarized light.
- Test targets need to be clean, free of lint, dust, smoke residue, fingerprints or other smudges, which, even when barely visible, can alter the measured colour of the patches.
- Reflective test targets should generally be non-glossy and transmissive test targets free of surface reflections (i.e. captured in a dark environment). Usable reflective targets include the 24 patch Macbeth ColourChecker™.

7.3 Capturing images of test targets

- The entire test target needs to be captured without obstruction of the target or the illumination.
- The test target needs to be lit with the lights, gels and power levels for which the scene analysis transform is desired. The resulting characterization is only valid for these spectral characteristics. Note that lights that measure at the same correlated colour temperature or chromaticity can have vastly different spectral power distributions and therefore result in different characterizations.
- The test target needs to be evenly illuminated. Reflection targets should be illuminated at 45 degrees from two opposing directions, except in daylight where the illumination of the sun at an angle of 45 degrees should be sufficient.
- To get a good representation of average camera flare, the test target should be surrounded by the average spectral power from the scene. It should not be placed on an unusual colour in the scene. For average scenes, the test target should be placed on a large neutral grey card, illuminated identically to the test target.
- The test target needs to be exposed to avoid clipping of any channels.
- The test target should be captured with the camera equipped with the optics that will be used in practice (lenses and filters), including any UV, IR, polarization, ND or colour conversion filters. The

exception is when a filter is used to intentionally alter the look of the resulting image, for example a warming filter, as otherwise the characterization will neutralize the effect of the filter.

- To avoid lens falloff and other off-centre camera behaviour, the test target should be centred in the camera field of view and cover 15 % to 30 % of the image width.
- For accurate capture of the test target colours, the image should include at least 200 pixels of each patch colour, excluding the edges of the patch.
- Evaluate the noise and evenness of an image of a test target as follows:
 - 1) Make a copy of the image of the test target
 - 2) Calculate the average image colour code value per patch excluding near the edges of the patch.
 - 3) Replace the actual code values per patch with the average code value per patch.
 - 4) Subtract the original image values, and amplify the difference image for evaluation.
 - 5) Examine the resulting image.
 - If the test target is unevenly lit, the resulting image will show colour or grey gradients in each patch.
 - If the test target is slightly stained, the stains will show up more clearly.
 - A well-captured test target will show only noise with no distinct marks or patterns.

8 Applying the scene analysis transform and encoding

Scene analysis transforms are applied as follows:

- a) Apply the inverse of the FP or Camera OECF to the camera values to produce linear response values (if necessary).
- b) Subtract the dark current (if necessary).
- c) Remove the estimated average camera flare (if desired).
 - The average camera flare can be removed by including a per-channel offset in the scene analysis transform. However the average camera flare is image specific. When processing images to scene-referred, it is better to separate out the camera flare subtraction which is image specific from the scene analysis transform which is adopted white specific.
 - Note that if the Camera OECF inverse was used for linearization, the flare subtracted is the difference between the flare present in the Camera OECF measurements and the estimated average flare for the particular image, and can be positive or negative for each channel.
- d) Apply the channel gains necessary to produce the selected adopted white to the linear response values to produce neutral-balanced linear response values.
 - If, after gain or OECF inversion the exposure range for any channel goes higher than that for the lowest channel, those channels should be clipped to the exposure value for the lowest channel.
 - This produces a neutral clipping ceiling, and in the case where camera clipping occurred on capture avoids artefacts resulting from the matrixing of clipped channel data with unclipped channel data.
 - Alternately, clipped channel information can be estimated from unclipped channel data using proprietary methods.
- e) Apply the scene analysis transform to the linear encoding colour space, based on the encoding primaries.

- f) Scale the linear colour space values to place scene highlights and mid-tones at reasonable levels.
 - This scaling is only preliminary, and should be re-evaluated later by viewing the image using the intended colour rendering transform and reference display.
- g) Apply any encoding colour component transfer function (not applicable for linear colour encodings).

9 Considerations for selecting scene analysis transforms

- The camera used and scene adopted white
- Expected spectral characteristics of scene colours to be captured
- Needs for accurate analysis of specific objects in the scene
- Scene analysis transform spectral limit produced
 - Different scene analysis transforms result in different spectral to colourimetric mappings
- Colour noise amplification

10 Suggested metadata for scene-referred colour encodings

10.1 Camera characteristics

- Camera gains (in the case of linear response) or OECFs (in the case of non-linear camera response)
- Camera dark current
- Camera flare average percent
- Camera spectral sensitivities
- Removable optical element spectral transmittances or standard designation

10.2 Colour encoding characteristics

- Colour encoding primaries (the nine relative tristimulus values)
- Colour component transfer function (if not linear, including any black or white scaling)
- Encoding adopted white chromaticity (for interpreting appearance)

10.3 Image specific characteristics

- Adopted white spectral power distribution (in the scene or at the focal plane, as applicable)
- Scene adopted white luminance
- Scene analysis transform applied, including any clipping or extrapolation of channels
- Scene analysis transform determination parameters
- Image state (scene- or focal-plane-referred; i.e. is camera flare included)