

How Creative Professionals Utilize Wide Color Gamut (WCG) and High Dynamic Range (HDR) Displays, and What Are the Applied Concerns?

Jack Holm

Tarkus Imaging, Carmel, CA
jack@tarkusimaging.com

Abstract

WCG/HDR content and displays are now ubiquitous, but have a wide variety of color capabilities. Systems supporting WCG/HDR images are necessarily more complex than standard dynamic range (SDR) systems, and this complexity is increased by WCG/HDR diversity. Practices must be rethought to accommodate. Solutions are at various stages of development in ITU and ISO. Content creators are on a learning curve. Additional considerations arise concerning viewer perceptions.

Author Keywords

wide color gamut (WCG), high dynamic range (HDR), color rendering, gamut mapping, color volume, color gamut, optical to optical transfer function (OOTF), opto-electronic transfer function (OETF), electro-optical transfer function (EOTF), camera controls, creative intent, look, color grading, single-master, reference white, diffuse white, HDR-focused, SDR-focused, dual-focused, observer metamerism, Farnsworth-Munsell 100 hue test, standard observer, golden eyes, viewing environment adaptation

1. Introduction

There is no question that WCG and HDR systems provide opportunities to dramatically improve user experiences in still and short burst photography, motion pictures and television, but these opportunities open the door to some potential problems:

- Colors outside the display color volume will be reproduced incorrectly. (The display color volume is the three-dimensional volume of colors the display can produce.)
- Increased display color variability due to metameric issues
- How to create and communicate multiple masters
- Co-optimization of masters with respect to a specific transform between them
- Adapting to other displays and viewing environments

The following sections will address these issues.

2. Color Rendering & Gamut Mapping

With WCG/HDR encodings, it is often possible to encode colors outside the mastering display color volume. Such colors can be found in natural scenes, and can be pushed outside by the color rendering as defined in ISO 22028-1:

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

In ITU-R BT terminology, the default color rendering is referred to as the reference Optical to Optical Transfer Function, or OOTF.

When displayed, colors encoded outside the mastering display color volume are mapped to some color in the display color volume, with the resulting color depending on the gamut mapping applied. While most gamut mapping algorithms attempt to choose a reasonable substitute, there can still be unexpected surprises and many-to-one mappings that hide artifacts that can be later revealed on displays with larger color volumes. Consequently, each color rendering of an image is specific to the display on which it was mastered, and ideally will not contain any encoded colors outside the mastering display color volume.

3. Effects of Color Rendering on Color Volume

Reference OOTFs: The combination of the ITU-R BT. 709 Opto-Electronic Transfer Function (OETF) with the ITU-R BT.1886 Electro-Optical Transfer Function (EOTF) results in a reference OOTF for SDR, which is also specified for the Perceptual Quantization (PQ) method specified in ITU-R BT 2100. Additionally, ITU-R BT.2100 specifies a different reference OOTF for the Hybrid Log-Gamma (HLG) method. The SDR OOTF provides a contrast and colorfulness boost in the midtone region, the latter being accomplished by applying the transfer function to the encoding Red, Green & Blue (RGB) channels. Because this OOTF is applied to the encoding RGB channels independently, it will not push any colors outside the encoding color volume.

Use of the SDR reference OOTF for PQ is complicated by the fact that in practice it is rarely used to map the input encoding range to the full PQ output range of 0 to 10,000 cd/m². The PQ encoding peak white was selected to support all HDR displays, without the intention that HDR displays should be able to produce the encoding peak white. Such high brightness can be beneficial for extremely bright transient events, or with very bright environments, but not generally. Also, in the case of very bright environments, content in the lower part of the range is not distinguishable. With PQ, different parts of the range are used at different times in line with the creative intent. This significantly increases what can be done creatively to affect the viewer experience, but also complicates cross-conversions to fixed range systems including SDR, and compensation for different end-user viewing environments (as will be discussed later).

The HLG reference OOTF is a linearly black and white scaled power function where the exponent (γ) depends on the peak white luminance (L_w) of the target display, as defined in Equation 1.

$$\gamma = 1.2 + 0.42 \text{Log}_{10} \left(\frac{L_w}{1000} \right) \quad (1)$$

Note that the value of γ will be > 1 so long as $L_w > 334 \text{ cd/m}^2$. Consequently, the HLG reference OOTF will increase contrast

in all tonal regions for displays with $L_W > 334$ (except possibly very near display black (L_B) when $L_B > 0$). It maps the encoding black and peak white to the display black and peak white. However, the HLG reference OOTF does not affect colorfulness, since it is applied in such a way as to preserve the ratios of the linear RGB channels. This is necessary because otherwise the colorfulness would vary between displays with different peak whites. It results in a somewhat less colorful “look” than with the SDR/PQ reference OOTF, but the colorfulness can be adjusted to be similar using camera controls or during grading. The ratio-preserving method of applying the reference HLG OOTF can push colors outside the encoding color volume, requiring the use of above peak white encoding values to fully address additive RGB display color volumes, which is a reason it is recommended to use narrow range encoding for HLG (e.g. 64-940 for 10-bits).

Figure 1 shows two SDR OOTFs in a log-log plot, along with three HLG OOTFs. Figure 2 shows how applying the SDR OOTF to ITU-R BT.2020 RGB channels independently affects image colors (including the channel clipping of colors outside the BT.2020 RGB triangle prior to applying the OOTF).




Figure 1. The ITU-R BT.709 – BT.1886 OOTF for a 0.1 to 100 cd/m² display (green), and a 0.005 to 1000 cd/m² display (black), along with HLG OOTFs for 0.005 to 1000 cd/m² (red), 0.0001 to 600 cd/m² (blue) and 0.01 to 4000 cd/m² (magenta) displays.

It is interesting to note the effects of the reference OOTFs on the input capture range and output dynamic range of images. For example, the SDR capture linear light range is 3942:1 (approximately 12 stops) for 10-bit code values 65 to 940, and the resulting linear light range for a 0.1 to 100 cd/m² display is 955:1 (approximately 10 stops). However, if a statistically average 160:1 pictorial dynamic range scene is exposed so the scene diffuse whites are placed one-half stop below the encoding peak white, the 160:1 scene range is expanded to 339:1 on the same display. Similarly, the reference HLG OOTF will expand a similarly placed 160:1 scene range to 434:1 on a 1000 cd/m² peak white display, and to 1558:1 on a 4000 cd/m² peak white display.

In contrast to the reference SDR OOTF, the reference HLG OOTF for 10-bit code values 65 to 940 has a much larger acceptance range of approximately 21 stops, 19 stops being below the HLG reference white. Compared to the reference SDR OOTF, the reference HLG OOTF accepts 7 stops lower and two stops higher in scene light, and maps the full scene acceptance range to the display luminance range.




Figure 2. The effects of applying the ITU-R BT.709 – BT.1886 OOTF with clipping in RGB. Red dots before OOTF, black dots after OOTF.

Creating HDR content for PQ: It is likely the reference SDR OOTF is specified for PQ in order to produce a default match between SDR displayed content and PQ displayed content, while allowing the content to be placed as desired within the PQ encoding luminance range. This thinking goes back to one of the fundamental reasons for HDR – to allow scenes with different brightness to be displayed at different brightness. While this can be done to a limited extent with SDR, since many scenes have a lower range than the reference SDR OOTF acceptance range, there is not that much extra room with SDR. For example, the 160:1 scene example above can be exposed one-half stop more, resulting in a displayed range from 0.25 to 100 cd/m² (395:1), but with absolutely no headroom for above white highlights. If it is exposed two stops less, the displayed range is 0.16 to 25 cd/m² (158:1).

On the other hand, if the output of the reference SDR OOTF is moved up and down the PQ display light range by multiplying it by different factors, the display light can be placed as desired, similarly to how the camera exposure is adjusted to place the image in the desired scene light range, while maintaining the default SDR “look.”

A problem arises, though, with scenes with inherent high dynamic ranges. These scenes are the other fundamental reason for HDR. Such scenes contain light sources (neutral or colored), specular reflections, or different areas with significantly different illumination levels. In this case simply scaling the reference SDR OOTF output higher can result in the linear part of the OETF being too high in the scene light range, degrading dark tones. This can be addressed by extending the gamma part of the OETF to higher scene light values instead of multiplying the display light by a factor.

In practice, camera controls are used to alter and shape the camera OETF used (and therefore the OOTF) on a scene-by-scene basis (as will be discussed later).

Creating HDR content with HLG: Since the reference HLG OOTF maps the full capture range to the full display range, it is analogous to the reference SDR OETF, with three caveats:

- It does not have a linear part
- It is not expected that the scene range will fill the reference HLG OOTF acceptance range or that the rendered image light will fill the display range
- It does not by default produce an increase in colorfulness

HLG is often used similarly to SDR, even to the extent of side-by-side shading/approval displays for live programs. In this case it is often desirable to match lightness of the SDR and HLG program material. This is done by setting the SDR display peak white to 203 cd/m² to match the nominal reference white of an HLG 1000 cd/m² peak white display. In this case, the program material is centered on the SDR dynamic range, but with the HLG content having about 2.3 stops of additional range at the top and 5.3 stops of additional range at the bottom. A transform/LUT is used to create the SDR stream from the HLG stream, where the transform manages highlights, shadows and colorfulness differently for the two streams as desired. This practice has the benefit of allowing optimization of the SDR content while still providing some added benefit with the HLG content.

The disadvantage of such workflows is the creative freedom to make the HDR content very light or very dark is reduced. For example, diffuse whites can't be placed above the nominal 203 cd/m² in the HDR content, because that would result in unacceptable clipping or highlight compression in the SDR stream. Also, while diffuse whites can be placed lower than 203 cd/m², to provide additional headroom or for creative effect, they can't be placed too much lower without negatively impacting the SDR stream. This is not altogether a bad thing, as it tends to keep the HLG content "out of the shadows," which is when creators go too dark with the content for creative purposes, not realizing they are making the content unwatchable in brighter viewing environments.

Occurrence of WCG/HDR Colors in Nature: When designing color encoding systems, consideration is given to which colors occur with some frequency in the natural world. Such colors are often called "surface colors" as they arise from the reflection of the illuminating source by a surface. This is a reasonable, but incomplete, consideration. There are many colors that are produced by refraction, interference, spectral sources, etc. that fall on or very near the spectral locus. Even typical surface colors are increased in purity if the light falling on them is of a similar color. A leaf by itself will measure as less saturated than when measured *in-situ* in a tree, because in the tree much of the illumination falling on the leaf is reflected from other leaves of the same color. Likewise, there are few surfaces that reflect less than 1% of the light falling on them, so if all that is needed is to capture surface colors accurately HDR is unnecessary.

Camera Adjustments and Grading: Other considerations in the need for and use of WCG and HDR relate to camera controls and grading. For example, the camera exposure controls where the top of the scene range falls in the capture encoding range and the camera black level setting controls where the bottom of the scene range falls. The black level setting can have a large effect on the apparent scene dynamic range. For

example, if the 160:1 scene mentioned above is exposed to place the scene whites one-sixth stop below the SDR encoding maximum, and the scene black level is set to one-sixth stop above the SDR encoding minimum, the 160:1 scene is expanded to 3129:1 in scene light (before the OOTF) resulting in a display light range of 835:1 (after the OOTF). Black level setting is one reason for the misconception that most natural scenes are inherently high dynamic range. The black level setting also can have a significant effect on colorfulness, as setting the scene black closer to encoding black effectively subtracts neutral light from all colors, resulting in increased saturation. While setting the black level is commonly done to align different cameras and make the images "look right," many camera operators and shaders do not realize that this simple operation has a profound effect on the actual OETF, even though the reference OETF may be applied after the black level adjustment has been made.

Other camera controls include:

1. Controls that effectively change the OETF tone curve:
 - "Gamma" controls that adjust lightness without moving the range endpoints by bending the OETF upwards or downwards
 - "Knee" controls that adjust the amount of highlight compression
 - "Black Gamma" controls that lighten or darken the darker colors without affecting the lighter colors
2. Controls that work in a HSV type color space:
 - "Saturation" controls how saturated the colors are
 - "Hue" provides the ability to rotate hues
 - "Saturated color lightness" allows saturated colors to be lightened or darkened without affecting neutrals
3. Menus that allow the camera operator to select from a list of OETFs and color modes providing different contrasts, dynamic ranges and color characteristics for different situations, either assuming the standard EOTF or that the encoded images will be further processed before output.

Image editing software and color grading stations provide even more flexibility for content creators to alter colors, including locally.

4. Metamerism

Metamerism is defined as the phenomenon where colors of two objects are perceived to be the same when viewed using a first illumination source, but are perceived to be different when viewed using a second illumination source with a different spectral power distribution. In practice, the word "metamerism" is used to describe other phenomena where two objects that have the same color in one situation have a different color in another situation.

The term "observer metamerism" refers to the case where two objects that appear the same to one observer appear different to a different observer when using the same illumination, resulting from the observers' eyes having different spectral sensitivities.

Observer metamerism is of particular importance with wide color gamut displays, because in order to produce wide color gamuts the displays must have primary emitters with relatively narrow spectral bandwidths. Narrow spectral bandwidths exacerbate observer metamerism because small differences in

the eye spectral sensitivities can result in significant differences in the visual signals.

Observer metamerism can be minimized by spreading out the spectral distributions of red, green and blue emitters in such a way as to keep the chromaticities of the emitters close to the spectral locus. The addition of other emitters, such as cyan or yellow, can also help reduce observer metamerism. However, adding emitters complicates display processing, since multiple possible combinations can produce the same color. They also change the shape of the addressable color volume, requiring different masters and/or re-rendering.

CIE Standard Observers: Even when displays are designed to minimize observer metamerism, unexpected colors can occur. In part, this can be attributed to differences between the CIE 2° and 10° Standard Observers and actual human observer spectral sensitivities. Research into better understanding of the HVS continues, but at present there is no perfect model of the HVS, so the only way to predict what someone will see is to have an actual person take a look.

Golden Eyes – Reference Human Observers: Observer metamerism is not new, and it has long been known that the CIE standard observers are not perfect. In fact, they can never be perfect due to human variability. The practical solution to this problem is to use specific human observers whose color vision is considered to be excellent.

Farnsworth-Munsell 100 Hue Test: Used to test color vision. Consists of 100 discs which are considered to have the same lightness and chroma but differ slightly in hue, forming a complete hue circle: red, yellow-red, yellow, green-yellow, green, blue-green, blue, purple-blue, purple, red-purple. Observers are asked to put the discs in hue order.

- Observers with “ideal” color vision will be able to “correctly” order the discs.
- Observers with low discrimination, including various forms of color blindness, will not be able to distinguish some hues, resulting in random errors in placement in at least some regions of the hue circle.
- Observers with high discrimination but anomalous color vision will place the discs in a different, but repeatable, order.

It is not always recognized that the 100 Hue Test tests not only color discrimination but also whether an individual’s color vision is consistent with the standard ordering of the hues. Color-critical roles in photography, graphic arts and motion pictures are typically filled by people who get perfect scores on this test, or have otherwise demonstrated that they have high discrimination *and* “ideal” color vision. In the motion picture industry, such people are sometimes called “golden eyes.”

Golden eyes can be expected to be reasonably consistent with each other in determining color matches, and can therefore be used to evaluate color in all cases, even where standard colorimetry based on CIE color matching functions does not work. Their judgment is accepted for color decisions, even though some other individuals may see things differently. They approve the look of the master, and then everyone else sees what they will see depending on their individual color vision.

One might ask, then, why not redo the color matching functions using golden eyes? The problem is that some aspects of how the golden eyes see color are a result of other factors than the

spectral response of the retina. A more complex and accurate model of the HVS is needed to achieve the color matching ability of golden eyes – improved color matching functions do not fully address the problem.

Since the professionals involved in making creative color choices and approving the master(s) are typically golden eyes, observer metamerism primarily causes problems when displays are used that have only been calibrated using measurements. Calibrations need to be verified visually, especially neutrals. Displays that are designed to minimize observer metamerism reduce the importance of visual verification but may not eliminate the need in critical situations.

5. Communicating Multiple Outputs

In the production of WCG/HDR content, especially content intended for diverse displays, it is essential to recognize three basic axioms:

1. The color appearance of content as specified by the master generally will not be the same as that of the scene captured, but will instead include enhancements to make it more pleasing to viewers, and look modifications to convey the desired intent of the creator.
2. The appearance enhancements and look modifications noted above will typically depend on the characteristics and capabilities of the display used for mastering the content.
3. The enhancements and look modifications made to produce the master appearance may need to be adjusted to reasonably maintain the creative intent on end-user displays with capabilities and viewing environments different from those of the master display.

Acceptance of these axioms is not guaranteed, as they require a more complex understanding of and approach to color reproduction. However, accepting them *ab initio* facilitates progress toward solutions based on a robust foundation.

Exactly how to approach content creation is also dependent on the use case. Still and short burst imaging systems can support complex color rendering in producing master images, and re-rendering for different output media can utilize sophisticated algorithms and extensive metadata. Motion picture and television feature imaging require support for complex color rendering, but re-rendering metadata options can be more limited. Color rendering options for live television are even more limited, often to camera controls and LUT boxes, requiring some sort of agreed re-rendering for HDR-SDR conversions.

Standards for communicating different master appearances, and adapting content for specific displays provide some options:

1. Communicating color rendered “scene light” imagery with a standardized method for adaptive rendering, *e.g.* the ITU-R BT.709/BT.2020 reference OETF followed by BT.1886 adaptive EOTF, or the BT.2100 HLG OETF & adaptive EOTF. In this case the imagery communicated is not scene-referred (estimated scene colorimetry) but includes the camera adjustments and look modifications that are the difference between the reference OETF and the actual OETF, the latter having been adjusted as needed to produce the desired color appearance on a standard reference display.
2. Communicating a single master HDR display-referred description with recommendations for adapting it to other

displays, e.g. ITU-R BT.2100 PQ EOTF & ITU-R BT.2408 EETF.

3. Communicating a single master HDR display-referred description with ancillary metadata specifying the characteristics of the master display (and viewing conditions), e.g. ISO 22028-5.
4. Communicating a master display-referred description with ancillary metadata specifying the characteristics of the master display (and viewing conditions), and metadata specifying the conversion to one or more other master displays e.g. ISO 21496-1.

In cases where there is metadata specifying the conversion, it can take different forms:

- Specification of a fixed tone curve or LUT that performs the transform
- Specification of an adaptive tone curve or LUT
- Specification of parameters or statistical values that can be used to derive or select adaptive tone curves or LUTs
- Specification of a gain map that maps between two different masters, i.e. an SDR master and an HDR master, or two different HDR masters

It should be noted that if only a tone curve is provided or derived, an application method is also required for color images. Some common application methods are defined in Annex 5 of ITU-R BT.2408:

1. $I_C T_C P$ – tone curve applied to I channel to create I_2 channel, $\text{Min}(I/I_2, I_2/I)$ multiplied by C_T, C_P to create C_{T2}, C_{P2}
2. $Y' C' B' C' R$ – tone curve applied to Y' channel, to create Y'_2 channel, $\text{Min}(Y'/Y'_2, Y'_2/Y')$ multiplied by C_B, C_R to create C_{B2}, C_{R2}
3. $YRGB$ – tone curve applied to Y' channel, ratio of linear Y values before and after tone curve then applied to linear R, G, B channels (used for HLG OOTF)
4. $R'G'B'$ – tone curve applied directly to $R'G'B'$ channels (used for SDR/PQ OOTF)
5. maxRGB – tone curve applied to whichever channel is highest, ratio of linear values before and after tone curve applied to other two channels

Each of these methods has distinct advantages and disadvantages, requiring creative decisions about which to use, as discussed in Annex 5 of ITU-R BT.2408. Different methods also can be blended as desired to produce intermediate results.

6. Multiple Masters and Co-Optimization

In cases where both SDR and HDR images and/or images color rendered for displays with different color primaries are needed, it may be desirable to co-optimize the content in order to facilitate the creation and communication of multiple versions. With still images, a gain map approach can be used to provide one master and a map to a different master of the same image. In this case there is no restriction on the optimization of each image. However, the degree to which intermediate images are useful will depend on whether they represent a likely third master display. For example, a PQ image with a gain map to SDR would not be useful for creating a third version for a RYRGB display since all three have different color volume shapes.

Another approach is to use an adaptive tone curve or LUT to define the transformation between the two versions. This approach allows for greater independence in the optimization of the versions, and the tone curve algorithm can be designed to produce intermediate renderings. This approach requires the specification of a tone curve application method for consistent results, and cannot address local differences.

When using the adaptive tone curve approach, the most important parameters are the differences between the endpoints (master display black and peak white) and the reference whites. The reference white is defined to be the output value corresponding to the peak white of a SDR master display. This value is unambiguous when one of the versions is in fact an SDR master, but can be ambiguous if no SDR master is available. It is also important to distinguish the reference white from what might be considered the “diffuse white.” While it is occasionally true that the reference white is set at what is assumed to be the diffuse white, in many cases they will be different. With SDR, the diffuse white often is set somewhat below the reference white so that a “knee” can be applied to compress above diffuse white highlights. In this case, the adopted diffuse white is probably around the “knee point” (where the knee starts).

Finally, if a fixed tone curve or LUT is used to transform between the versions, co-optimization (including aligning the reference whites) will be necessary.

7. Viewing Environment Compensation

In most cases, masters are prepared for viewing in specified, standardized viewing environments that are relatively dim. Dark content that is easily visible in such environments may be almost impossible to see in bright viewing environments, such as in rooms with windows and no heavy curtains in the daytime, or even outdoors. One area where display manufacturers could do a great deal to improve the user experience is by including or improving display adaptations for the viewing environment. In the meantime, creative professionals need to be mindful of this issue when releasing non-theatrical versions.

8. Summary

There are many reasons creative professionals value WCG and HDR capabilities, but the reasons often have more to do with the color rendering and creative intent than the original scene color gamut and dynamic range.

When performing creative adjustments, care should be taken to avoid pushing colors outside the mastering display color volume. As an added safety measure, the content may be clipped to the mastering display color volume before distribution.

Display calibrations should be verified visually using people with known high discrimination and normal color vision, especially the neutral scale.

Consideration should be given to the degree to which co-optimization of masters is desirable, and to the nature of the transforms used to go between different masters. There are many transform options and no single correct answer.

A current issue that results in some user dissatisfaction is very dark HDR content and the inability of most HDR displays to adjust the content for different user viewing environments.

9. References

ISO and ITU-R standards as cited in text.

1. ISO 22028-1:2016(E), Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange — Part 1: Architecture and requirements; Second Edition; 2016
2. ISO/DTS 22028-5:2023(E), Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange — Part 5: High dynamic range and wide colour gamut encoding for still images (HDR/ WCG)
3. ISO/DIS 21496-1:2024(E), Digital photography — Gain map metadata for image conversion — Part 1: Dynamic range conversion
4. ITU-R BT.709-6, Parameter values for the HDTV standards for production and international programme exchange (2015)
5. ITU-R BT.1886, Reference electro-optical transfer function for flat panel displays used in HDTV studio production (2011)
6. ITU-R BT.2020-2, Parameter values for ultra-high definition television systems for production and international programme exchange (2015)
7. ITU-R BT.2100-2, Image parameter values for high dynamic range television for use in production and international programme exchange (2018)
8. ITU-R BT.2408-5, Guidance for operational practices in HDR television production (2022)