The HDR Fact Book

Milano HDR Live Workshop Edition

With compliments from

Video Progetti
A feeling for technology.

Contents

EDITOR'S NOTE

This publication has been created by collecting publicly available information and articles. It is intended for educational purposes only and it is not for sale.

The editor wishes to thank and credit the following institutions and their authors for their contribution to the materials contained herein:

acescolorspace.com
Bruce Devlin aka Mr. MXF
DisplayDaily.com
Dolby
Unravel Mastering Services
Wikipedia (through editorial supervision)

Other contributions are credited as part of the articles or images.

ACES Color Space

The Academy Color Encoding System (ACES) is the industry standard for managing color throughout the life cycle of a motion picture or television production. From image capture through editing, VFX, mastering, public presentation, archiving and future remastering, ACES ensures a consistent color experience that preserves the filmmaker's creative vision.

ACES is a free, open, device-independent color management and image interchange system that can be applied to almost any current or future workflow.

The most current release version, ACES 1.2, clarifies some specifications and addresses a few user-generated comments and requests and is being widely used around the globe. ACES 2.0, currently in development, will simplify and add support for additional workflows.

ACES includes the following components:

- a collection of wide-gamut colorspaces (ACES2065-1, ACEScg, ACEScc, & ACEScct).
    - based on the RGB color model.
    - image data is scene-referred, i.e. the numerical values are related to the original scene lighting, as reflected or emitted from the real objects & lights on the set at the time of filming.
- transforms to convert images captured on specific hardware to an ACES colorspace (Display-Referred to Scene-Referred)
- transforms to convert images in other popular colorspaces to an ACES colorspaces (Display-Referred to Scene-Referred)
- a number of utility transforms to help convert images for specific uses in a pipeline (Utility - sRGB - texture to ACEScg for use on a color channel on a cg material)
- transforms which allow you to convert an ACES image to

another colorspace (ACEScg to sRGB, ACEScc to Rec. 709).

- transforms which allow an artist to properly view ACES images on calibrated hardware (such as on computer monitors and televisions).

AV-1

AOMedia Video 1 (AV1) is an open, royalty-free video coding format initially designed for video streaming over the Internet. It was developed as a successor to VP9 by the Alliance for Open Media (AOMedia)*, a consortium founded in 2015 that includes semiconductor firms, video on demand providers, video content producers, software development companies and web browser vendors. The AV1 bitstream specification includes a reference video codec. In 2018, Facebook conducted testing that approximated real-world conditions, and the AV1 reference encoder achieved 34%, 46.2% and 50.3% higher data compression than libvpx-vp9, x264 High profile, and x264 Main profile respectively.

Like VP9, but unlike H.264/AVC and HEVC, AV1 has a royalty-free licensing model that does not hinder adoption in open-source projects.

* The governing members of the Alliance for Open Media are Amazon, Apple, ARM, Cisco, Facebook, Google, Huawei, Intel, Microsoft, Mozilla, Netflix, Nvidia, Samsung Electronics and Tencent.

AVC

Advanced Video Coding (AVC), also referred to as H.264 or MPEG-4 Part 10, is a video compression standard based on block-oriented, motion-compensated coding. It is by far the most commonly used format for the recording, compression, and distribution of video content, used by 91% of video industry

developers as of September 2019. It supports a maximum resolution of 8K UHD.
It is also widely used by streaming internet sources, such as videos from Vimeo, YouTube, and the iTunes Store, web software such as the Adobe Flash Player and Microsoft Silverlight, and also various HDTV broadcasts over terrestrial (ATSC, ISDB-T, DVB-T or DVB-T2), cable (DVB-C), and satellite (DVB-S and DVB-S2).

In terms of its origin, AVC was developed jointly by the International Telecommunications Union (ITU-T) and the Moving Picture Experts Group (MPEG), which is a project of the ISO/IEC and familiar to many users because of popular and accessible MPEG file formats, like .mpg.

It is not royalty free and currently 20 organizations are patent holders within MPEG LA patent pool for AVC.

Barten Curve / Ramp

The Barten Ramp is an extremely handy graph that plots out where most people can begin to see banding in a gradient (i.e. the steps between each shade) when mapping out all the way to 10,000 nits for potential HDR imagery. The area in green (see next page diagrams) shows where no banding can be seen but the area in red shows where banding can be seen and is therefore problematic.
For image quality, we would prefer to have an image stay below the Barten threshold at all times. As can be seen however, in order to do that with a standard gamma transfer function, we would need to allow for 15-bits per channel. If we only allocated 10-bits per channel, our darker regions would still exhibit banding artifacts. In addition to 15-bits per channel being a very high requirement for commercial systems (and even more so for

Barten Ramp and PQ

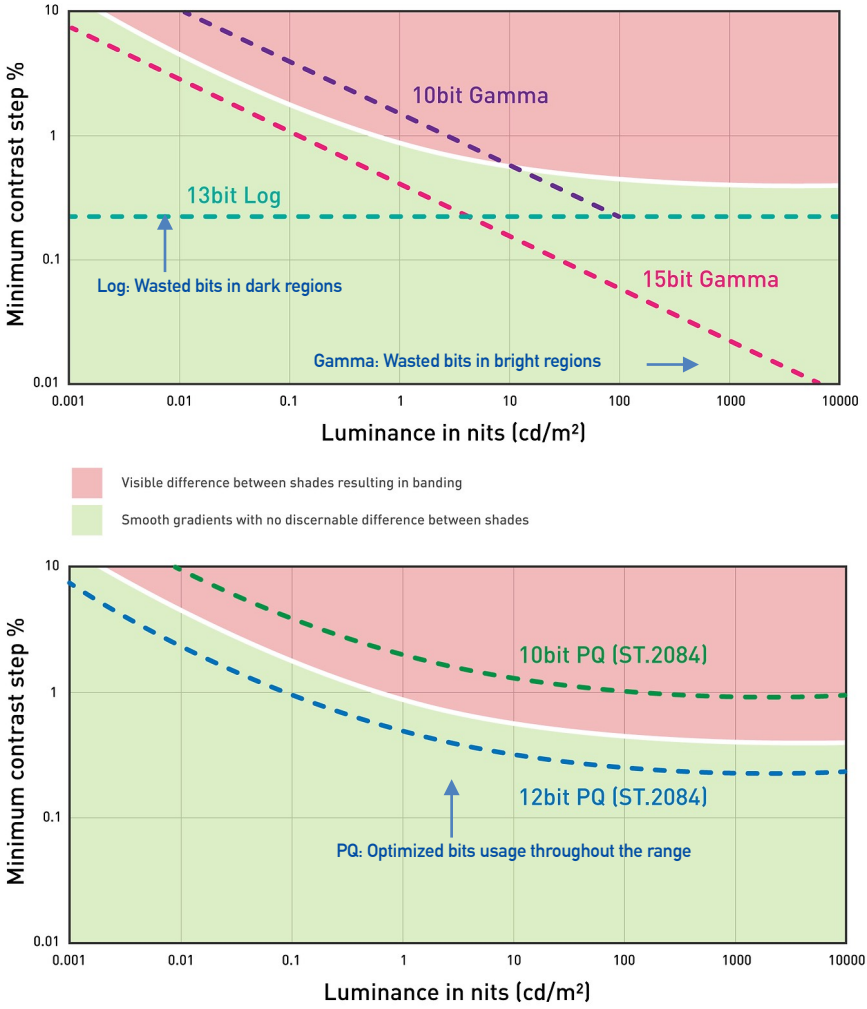

consumer systems!), it's also a very inefficient use of the data; there's far more data assigned in the mid and bright portions of the image than what is required. A 13-bit Log encoded image stays below the Barten threshold also but has an inefficient encoding of the shadows where more values are allocated than required.
What we really want is a curve that stays below the Barten threshold and follows the curve as close as possible to maintain optimal efficiency in encoding. Solution for this is the "PQ" transfer function (ST. 2084) (See page 37).

(Source: https://www.unravel.com.au/understanding-gamma)

Bit Depth or Colour Depth

Color depth, also known as bit depth, is either the number of bits used to indicate the color of a single pixel, or the number of bits used for each color component of a single pixel.

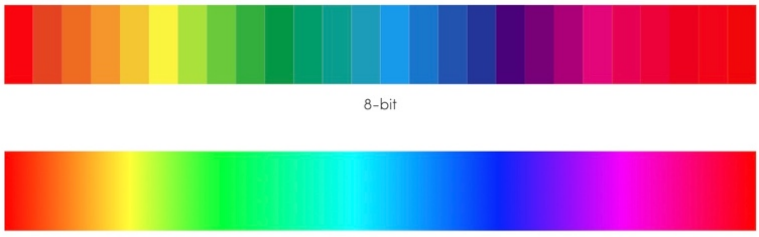

Credits: B&H Photo Video
Editor's note: scaling on 8 bit example exaggerated for illustration purpose only

Color depth is only one aspect of color representation, expressing the precision with which the amount of each primary can be expressed; the other aspect is how broad a range of colors can be expressed (the gamut). The definition of both color

precision and gamut is accomplished with a color encoding specification which assigns a digital code value to a location in a color space. The colour depth for HD content is typically 8 bits (10 bits for mastering), Ultra HD and High Dynamic Range (HDR) content is typically between 10 and 12 bits for distribution and up to 16 bits for mastering. (Editor's note: for the rationale behind requirements for higher bit depth in connection with HDR see also Barten Ramp, page 9).

Bitrate

In telecommunications and computing, bit rate is the number of bits that are conveyed or processed per unit of time.
The bit rate is expressed in the unit bit per second (symbol: bit/s), often in conjunction with an SI prefix such as kilo (1 kbit/s = 1,000 bit/s), mega (1 Mbit/s = 1,000 kbit/s), giga (1 Gbit/s = 1,000 Mbit/s) or tera (1 Tbit/s = 1,000 Gbit/s).
The non-standard abbreviation bps is often used to replace the standard symbol bit/s, so that, for example, 1 Mbps is used to mean one million bits per second.
In most computing and digital communication environments, one byte per second (symbol: B/s) corresponds to 8 bit/s.

Bitrate has an association with the quality of an audio or video file. A 3 Mbps may look better than the same file compressed at 1 Mbps using same codec.

Color Calibration

The aim of color calibration is to measure and/or adjust the color response of a device (input or output) to a known state. In

International Color Consortium (ICC) terms, this is the basis for an additional color characterization of the device and later profiling. In non-ICC workflows, calibration refers sometimes to establishing a known relationship to a standard color space in one go. The device that is to be calibrated is sometimes known as a calibration source; the color space that serves as a standard is sometimes known as a calibration target.
Color calibration is a requirement for all devices taking an active part of a color-managed workflow, and is used by many industries, such as television production, gaming, photography, engineering, chemistry, medicine and more.

Color Spaces

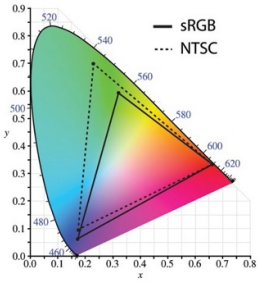

A color space is a specific organization of colors. In combination with color profiling supported by various physical devices, it supports reproducible representations of color – whether such representation entails an analog or a digital representation. A color space may be arbitrary, i.e. with physically realized colors assigned to a set of physical color swatches with corresponding assigned color names (including discrete numbers in – for example – the Pantone collection), or structured with mathematical rigor (as with the NCS System, Adobe RGB and sRGB). A "color space" is a useful conceptual tool for understanding the color capabilities of a particular device or digital file.

In cinema and television RGB representation through values associated to primaries (Red, Green, and Blue) or YUV (components) is mainly used.

Other color spaces such as XYZ and Lab are more representative of the human color vision model.

Contrast Ratio

The contrast ratio (CR) is a property of a display system, defined as the ratio of the luminance of the brightest shade (white) to that of the darkest shade (black) that the system is capable of producing. A high contrast ratio is a desired aspect of any display. It is represented as a ratio of n:1.

CRI (Color Remapping Information)

A set of standardized metadata which can be used in mapping one color space to another. This type of information can be useful when converting from Wide Color Gamut (WCG) video to another format or from HDR to SDR for example. Color Remapping Information was adopted in ISO/IEC 23008-2:2014/ ITU-T H.265:2014 High Efficiency Video Coding (HEVC) specification and is being implemented in the Ultra HD Blu-ray specification. It is also being considered in SMPTE ST 2094.

(Source: Google Patents)

DCI-P3, D65-P3, ST 428-1

A digital cinema color space. DCI-P3 is an RGB color space first defined in 2005 as part of the Digital Cinema Initiative, to be used for digital theatrical motion picture distribution. Display P3 is a variant developed by Apple Inc. for wide-gamut displays.
It has been standardized in 2006 by SMPTE ST 428-1.
This color space features a Color Gamut that is much wider than sRGB (see Rec. 709 p. 36).

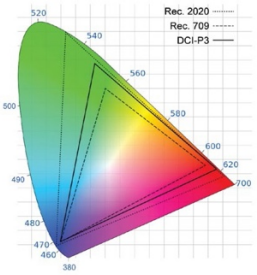

All Digital Cinema Projectors are capable of displaying the DCI-P3 color space in its entirety. D65-P3 means that the color

temperature of the white point is set at D65 instead of the “DCI” white point.

The illustration shows: The large color space proposed by Rec. 2020, the new standard for Ultra HD TVs, (only currently fully achievable on laser projectors). The smaller DCI-P3 color space (Digital Cinema) and the smallest Rec. 709 space (traditional video monitors, including HD - Broadcast TV, Blu-ray, Over-The-Top).

Display Referred vs. Scene Referred

The two terms are often used in association with HDR workflows.
As a general rule, in a live TV environment, you want to capture the scene in full HDR and send a signal that's about right for every device without any specialist setup for the brightness / gamut of that devices. The process is known as scene referred (and it is at the base of HLG).
For movies and scripted content, you want to optimize the picture so that every display is running at maximum quality. This is display referred and is why the properties of the display affect the signal that gets given to the pixels. Display Referred is at the base of all PQ workflows (e.g. HDR10, HDR10+ and Dolby Vision).
Display referred often has more metadata and system complexity, but the results can be better when tuned. Scene referred is often simpler and less likely for a consumer to break through attempting optimization.

The difference between “scene referred” and “display referred” can be also explained according to where the rendering intent is implemented. In scene referred for example (e.g. HLG), the

rendering intent is implemented in the EOTF at the display after the application of the inverse OETF. The rendering intent stage is needed to account for the psychovisual effects of viewing a display in a dim environment.

In display referred content (e.g. PQ), the rendering intent is decided at the camera, then the inverse EOTF is applied. At the display, the EOTF then undoes the non-linearity introduced by the inverse EOTF and recreates a perceptually accurate image. In theory, both should work identically and produce the same non-linear system gamma or rendering intent. These are shown schematically below.

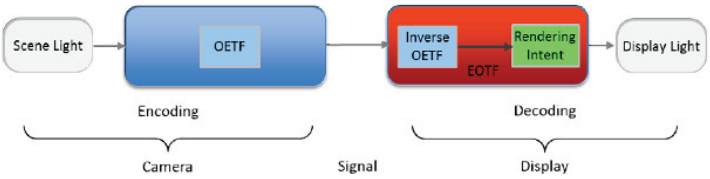

Case 1 : Scene Referred

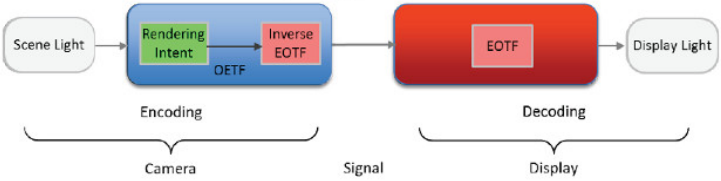

Case 2 : Display Referred

Source: DisplayDaily.com

Dynamic Range

Dynamic range is the ratio between the largest and smallest values that a certain quantity can assume. It is often used in the context of signals, like sound and light. It is measured either as a ratio or as a base-10 (decibel) or base-2 (doublings, bits or stops) logarithmic value of the difference between the smallest and largest signal values.

Dolby Vision, Dolby Vision IQ

Dolby Vision is a set of technologies developed by Dolby Laboratories for high dynamic range (HDR) video.
It covers content creation, distribution, and playback. It includes dynamic metadata that are used to adjust and optimize each frame of the HDR video to the consumer display's capabilities in a way based on the content creator's intents.
Dolby Vision was introduced in 2014, making it the first available HDR format. HDR10+ is a competitor HDR format that also uses dynamic metadata. (See page 23)
Dolby Vision IQ is an update designed to optimize Dolby Vision content according to the ambient light.

Other facts:
- EOTF:PQ
- Color gamut: Rec. 2020
- Pixel depth: 10 or 12-bit
- Max luminance: Although Dolby Vision standard technically supports up to 10,000 nits, content is mastered usually at 1,000 nits.
- Metadata: Dynamic, can adjust Tone Mapping at the display on a per-scene or frame basis.

EDID

Extended Display Identification Data (EDID) and Enhanced EDID (E-EDID) are metadata formats for display devices to describe their capabilities to a video source (e.g., graphics card or set-top box). The data format is defined by a standard published by the Video Electronics Standards Association (VESA).
The EDID data structure includes manufacturer name and serial number, product type, phosphor or filter type (as chromaticity data), timings supported by the display, display size, luminance data and (for digital displays only) pixel mapping data.
DisplayID is a VESA standard targeted to replace EDID and E-EDID extensions with a uniform format suited for both PC monitor and consumer electronics devices.

EOTF

The electro-optical transfer function (EOTF) is the transfer function having the picture or video signal as input and converting it into the linear light output of the display. In other terms it is a mathematical function that maps digital code values to displayed luminance. The process is performed within the display device.

See also OETF, ST 2084 (page 32 and page 37).

Flicker

Flicker is a visible change in brightness between cycles displayed on video displays. It applies to the refresh interval on cathode ray tube (CRT) televisions and computer monitors, as well as plasma computer displays and televisions. It results in an

undesirable changing of brightness mainly visible at frequencies below 50 frames per seconds. Flicker perception by human eye is dependent on the displayed picture brightness displays. For this reason high luminance displays may require higher frame rates.

Frame Rate

Frame rate (expressed in frames per second or FPS) is typically the frequency (rate) at which consecutive images (frames) are captured or displayed. This definition applies to film and video cameras, computer animation, and motion capture systems. In these contexts, frame rate may be used interchangeably with frame frequency and refresh rate, which are expressed in hertz. Additionally, in the context of computer graphics performance, FPS is the rate at which a system, particularly a GPU, is able to generate frames, and refresh rate is the frequency at which a display shows completed frames. In electronic camera specifications frame rate refers to the maximum possible rate frames could be captured, but in practice, other settings (such as exposure time) may reduce the actual frequency to a lower number than the frame rate.

In film, television, or video, frame rate is a critical factor for synchronizing audio with pictures.

f-stop (and Dynamic Range)

An f-stop or f-number is a measure of the light-gathering ability of an optical system such as a camera lens. It is calculated by dividing the system's focal length by the diameter of the entrance pupil (“clear aperture”). The f-number is a key factor in

determining the depth of field, diffraction, and exposure of an image.
The f-stop is often used to express the contrast ratio within a picture, more correctly referred to with the photographic term *latitude*.
As an example: a camera capable of handling 10 f-stops, means that the contrast value (ratio between white and black) can reach 1024:1 (in fact 2^{10} =1024).
As a reference, the human eye can handle 20 f-stops whereas SDR video images 6 -7 stops.

Gamut or Color Gamut

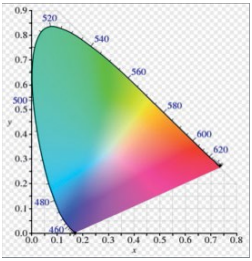

In color reproduction (including computer graphics and photography), the gamut, or color gamut, is a certain complete subset of colors. The most common usage refers to the subset of colors which can be accurately represented in a given circumstance, such as within a given color space or by a certain output device.
Another sense, less frequently used but still correct, refers to the complete set of colors found within an image at a given time. In this context, digitizing a photograph, converting a digitized image to a different color space, or outputting it to a given medium using a certain output device generally alters its gamut, in the sense that some of the colors in the original are lost in the process.

Gamut Mapping

In nearly every translation process (that is the transformation of the representation of a color from one color space to another), we have to deal with the fact that the color gamut of different

devices vary in range which makes an accurate reproduction impossible.

These therefore need some rearrangement near the borders of the gamut. Some colors must be shifted to the inside of the gamut, as these otherwise cannot be represented on the output device and would simply be clipped. This so-called gamut mismatch occurs for example, when we translate from the RGB color space with a wider gamut into the CMYK color space with a narrower gamut range.

The color management system can utilize various methods to achieve desired results and give experienced users control of the gamut mapping behavior.

HDMI

Stands for High-Definition Multimedia Interface (HDMI), a proprietary audio/video interface for transmitting uncompressed video data and compressed or uncompressed digital audio data from an HDMI-compliant source device, such as a display controller, to a compatible computer monitor, video projector, digital television, or digital audio device.

HDMI implements the EIA/CEA-861 standards, which define video formats and waveforms, transport of compressed and uncompressed LPCM audio, auxiliary data, and implementations of the VESA EDID. The Consumer Electronics Control (CEC) capability allows HDMI devices to control each other when necessary and allows the user to operate multiple devices with one handheld remote control device.

Several versions of HDMI have been developed and deployed since the initial release of the technology, occasionally introducing new connectors with smaller form factors, but all versions still use the same basic pinout and are compatible with

all connector types and cables. Other than improved audio and video capacity, performance, resolution and color spaces, newer versions have optional advanced features such as 3D, Ethernet data connection, and CEC extensions.

HDR

High dynamic range imaging (HDRI) refers to the set of imaging technologies and techniques that allow to increase the dynamic range of images or videos. It covers the acquisition, creation, storage, distribution and display of images and videos. This can be used to portray more realistic images with higher contrast, darker darks and brighter brights.

HDR10

HDR10 is an open high-dynamic-range video (HDR) standard announced on 27 August 2015 by the Consumer Technology Association. It is the most widespread of the HDR formats.
HDR10 is not backward compatible with SDR. It includes HDR static metadata but not dynamic metadata. It doesn't offer the ability to optimize content to the consumer display's capabilities in a way based on the content creator's intent.

Other facts:

- EOTF: SMPTE ST 2084 (PQ)
- Bit depth: 10 bit
- Color primaries: ITU-R BT.2020 (identical to BT.2100 primaries)
- Static metadata: SMPTE ST 2086 (mastering display color volume), MaxFALL (maximum frame-average light level), and MaxCLL (maximum content light level)

- Color sub-sampling: 4:2:0 (for compressed video sources)

HDR10 is technically limited to a maximum of 10,000 nits peak brightness, however common HDR10 contents are mastered with peak brightness from 1,000 to 4,000 nits.

HDR10+

HDR10+, also known as HDR10 Plus, was announced on 20 April 2017, by Samsung and Amazon Video. HDR10+ updates HDR10 by adding dynamic metadata that can be used to more accurately adjust brightness levels up to the full range of PQ code values (10,000 nits maximum brightness) on a scene-by-scene or frame-by-frame basis. The technology is standardized and defined in SMPTE ST 2094-40.
HDR10+ is an open standard and is royalty-free; it is supported by a growing list of post-production software and tools.
A certification and logo program for HDR10+ device manufacturers is available with an annual administration fee for certain adopter categories and no per-unit royalty.

HDR10+ signals the dynamic range and scene characteristics on a scene-by-scene or even frame-by-frame basis. The display device then uses the dynamic metadata to apply an appropriate tone map through the process of dynamic tone mapping.

Compatibility: Backward compatible for viewing on HDR10 displays. HDR10+ and Dolby Vision do not use the same dynamic metadata and are therefore not compatible.

Other facts:
- EOTF: SMPTE ST 2084 (PQ)
- Chroma subsampling: 4:2:0 (for compressed video

sources)

- Resolution: Agnostic
- Bit depth: 10-bit or more (up to 16-bit) per color channel
- Color primaries: ITU-R BT.2020
- Maximum linearized pixel value: 10,000 cd/m2 for each color R/G/B (content)
- Metadata (required): Mastering Display Color Volume Metadata[24]
- Metadata (optional): MaxCLL, MaxFALL[25]
- HDR10+ technology can support the full range of HDR standards to 10,000 Nits, 8K and BT.2020 color gamut. Being resolution agnostic, metadata needs to be created only once and can be applied to any target resolution.
- HDR10+ is supported in video encoding technologies HEVC, AV1, VP9 compatibility via WebM as well as any codec that supports ITU-T T.35 metadata.

HEVC

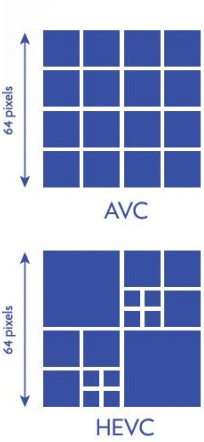

High Efficiency Video Coding (HEVC), also known as H.265 and MPEG-H Part 2, is a video compression standard designed as part of the MPEG-H project as a successor to the widely used Advanced Video Coding (AVC, H.264, or MPEG-4 Part 10). In comparison to AVC, HEVC offers from 25% to 50% better data compression at the same level of video quality, or substantially improved video quality at the same bit rate. It supports resolutions up to 8192×4320, including 8K UHD, and unlike the primarily 8-bit AVC, HEVC's higher fidelity Main 10 profile has been incorporated into nearly all supporting hardware.

While AVC uses the integer discrete cosine transform (DCT) with 4×4 and 8×8 block sizes, HEVC uses both integer DCT and discrete sine transform (DST) with varied block sizes between 4×4 and 32×32. The High Efficiency Image Format (HEIF) is

based on HEVC.

High Frame Rate

In motion picture technology—either film or video—high frame rate (HFR) refers to higher frame rates than typical prior practice. The frame rate for motion picture film cameras was typically 24 frames per second (fps) with multiple flashes on each frame during projection to prevent flicker.

When considering HFR as part of UHDTV, two high frame rates have been specified: 100 Hz (for 50 Hz countries) and 120 Hz (for 60 Hz countries).

HLG

The hybrid log–gamma (HLG) transfer function is a transfer function jointly developed by the BBC and NHK for high dynamic range (HDR) display.

It is backward compatible for viewing on SDR displays, provided these are able to interpret Rec.2020 colorimetry.

It was approved as ARIB STD-B67 by the Association of Radio Industries and Businesses (ARIB). It is also defined in ATSC 3.0, Digital Video Broadcasting (DVB) UHD-1 Phase 2, and International Telecommunication Union (ITU) Rec. 2100.

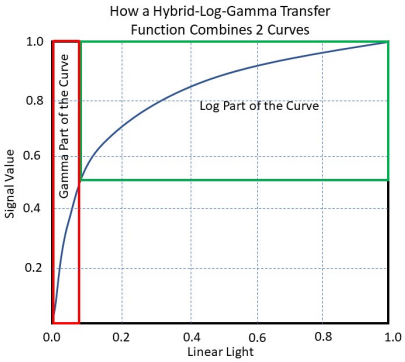

Credits: Telestream Blog, Steve Bilow

HLG extends the conventional SDR gamma curve by adding a logarithmic curve for the upper half of the signal values, which allows for a larger dynamic range.

HLG is an HDR format that uses the HLG transfer function, BT.2020 color primaries and a bitdepth of 10-bit.

Both HLG transfer function and the HLG format are royalty-free

HLG is used by video services such as BBC iPlayer, DirecTV, Freeview Play, and YouTube.

Other facts:
- Release: 2015 by BBC and NHK
- Royalty-free open standard
- EOTF:HLG
- Color gamut: Rec. 2020
- Pixel depth: 10-bit
- Max luminance: 1,000 nits
- Metadata:None
- Compatibility: Backward compatible for viewing on SDR displays, provided these are able to interpret Rec.2020 colorimetry.

Image Resolution

Image resolution is the detail an image holds. The term applies to digital images, film images, and other types of images.
Image resolution can be measured in various ways. Resolution quantifies how close lines can be to each other and still be visibly resolved. Resolution units can be tied to physical sizes (e.g. lines per mm, lines per inch), to the overall size of a picture (lines per picture height, also known simply as lines, TV lines, or TVL), or to angular subtense. Instead of single lines, line pairs are often used, composed of a dark line and an adjacent light line; for example, a resolution of 10 lines per millimeter means 5 dark lines alternating with 5 light lines, or 5 line pairs per millimeter (5

LP/mm). Photographic lens and are most often quoted in line pairs per millimeter.

In digital television it is common to use the number of pixel count per every dimension of the 2D picture space.

For instance, “HD” is implying 1920 x 1080 resolution (with 16:9 1,77:1 aspect ratio). Ultra High Definition (Ultra HD) has a resolution of 3840 x 2160 pixels, will display accurately on 16:9 aspect ratio (1.77:1).
“4K” digital cinema projectors have a resolution of 4096 x 2160 pixels although most theatrical cinema content is projected at either 4096 x 1716 (2.39 aspect ratio) or 3996 x 2160 (1.85 aspect ratio).

IMF

Interoperable Master Format (IMF) is a container format for the standardized digital delivery and storage of finished audio-visual masters, including movies, episodic content and advertisements.
In IMF, each kind of essence (video, audio, subtitles, etc.) is stored in individual media files (“Track Files”) and the instructions for synchronizing these media files are stored in a separate playlist file (“Composition Playlist”). This component-based approach allows the large media files to be reused across multiple playlists, each representing a different version of the content.
The IMF family of standards, ST 2067, is maintained by SMPTE. Its first edition was published in 2013.
IMF is closely related to Digital Cinema Package (“DCP”), leveraging several of the same standards for packaging and architecture.

Inverse Tone Mapping (ITM)

Inverse tone mapping is the inverse technique with respect to Tone Mapping (see page 39) that allows to expand the luminance range, mapping a low dynamic range image into a higher dynamic range image. It is notably used to upscale SDR videos to HDR videos.

LCD

A liquid-crystal display (LCD) is a flat-panel display or other electronically modulated optical device that uses the light-modulating properties of liquid crystals combined with polarizers. Liquid crystals do not emit light directly but instead use a backlight or reflector to produce images in color or monochrome.

LCDs are used in a wide range of applications, including LCD televisions, computer monitors, instrument panels, aircraft cockpit displays, and indoor and outdoor signage. Small LCD screens are common in LCD projectors and portable consumer devices such as digital cameras, watches, calculators, and mobile telephones, including smartphones. LCD screens have replaced heavy, bulky and less energy-efficient cathode-ray tube (CRT) displays in nearly all applications. The phosphors used in CRTs make them vulnerable to image burn-in when a static image is displayed on a screen for a long time, e.g., the table frame for an airline flight schedule on an indoor sign. LCDs do not have this weakness, but are still susceptible to image persistence.

Common technologies used for LCD's are twisted nematic (TN) (mainly used for monochrome screens, e.g. watches), IPS (in-

plane switching) and VA (vertical alignment).

LCEVC

Low Complexity Enhancement Video Coding (LCEVC) is a ISO/IEC video coding standard developed by the Moving Picture Experts Group (MPEG) under the project name MPEG-5 Part 2 LCEVC.
LCEVC specifies an enhancement layer which, when combined with a base video encoded with a separate codec, produces an enhanced video stream. The enhancement layer provides improved features to existing codecs, such as compression capability extension and lower encoding/decoding complexity, for live streaming or broadcasting applications.
LCEVC leverages a base video codec (e.g., AVC, HEVC, VP9, AV1, EVC or VVC) and employs an efficient low-complexity enhancement that adds up to two layers of encoded residuals, along with normative signalled up-sampling methods, that correct artifacts produced by the base video codec and add detail and sharpness for the final output video.

It provides additional compression efficiency to any existing or future video codec and reduces the processing complexity of encoding and decoding.
LCEVC can be implemented with software updates for encoders and decoders, and was designed to leverage available hardware acceleration for graphics processing.

Licensing is free of cost per decoder or encoder, but broadcaster or encoder-side use of enhancement layers is subject to fees.
LCEVC is proprietary to V-Nova and subject to V-Nova's proprietary license.

LUT

In digital imaging, LUT stands for Look-Up Table. A Lookup table or LUT is an array of numbers with values for color input and output. For every color value processed with the table, the LUT will transform or shift it to a different value. In simple terms, it is a preset for altering the color look of cinematography or photography with post-production software. Color grading with LUTs can be technical and transform footage from one color space to another, or creative and provide a desired or cinematic look.

A 1D LUT or a 1D curve, operates on a single dimension of color or intensity. It takes input values along one axis and provides corresponding output values along another axis. For example, in a typical color grading application, a 1D LUT might take input values for red, green, and blue channels and produce new output values for those channels.

3D LUT allows for more complex color grading as it controls color and luminance in a three-dimensional color space through hue, saturation, and brightness. The LUT files for 3D LUTs typically have the extension ".3dl" or '.cube' to distinguish them from one-dimensional lookup tables. 3D LUTs are often used in high-end color grading and calibration processes where accurate color reproduction is crucial.

MaxCLL Metadata

Maximum Content Light Level (MaxCLL) is an integer metadata value defining the maximum light level of any single pixel (in cd/m2 or nits) of the entire playback sequence. MaxCLL is usually measured off the final delivered content after mastering. If one uses the full light level of the HDR mastering display and adds a hard clip at its maximum value, MaxCLL would be equal to the

peak luminance of the mastering monitor.

(Source: https://professionalsupport.dolby.com/s/article/Calculation-of-MaxFALL-and-MaxCLL-metadata?language=en_US)

MaxFALL Metadata

Maximum Frame Average Light Level (MaxFALL) is an integer metadata value that indicates the maximum value of the frame average light level (in cd/m2 or nits) of the entire playback sequence. MaxFALL is calculated by averaging the decoded luminance values of all the pixels within a frame. MaxFALL is usually much lower than MaxCLL.

(Source: https://professionalsupport.dolby.com/s/article/Calculation-of-MaxFALL-and-MaxCLL-metadata?language=en_US)

MiniLED

Mini LED is an improved version of the conventional LCD backlight technology. With the sole purpose to selectively light up the display, mini LEDs are less than 0.2mm and can offer better control of the backlighting which leads to better image quality. A device like a TV consists of an LCD panel with LEDs for backlighting and these Mini LEDs, arranged in array, are lit up completely or dimmed down for darker scenes.

Some common advantages of mini LEDs over OLEDs are:

- Less manufacturing cost
- Better at displaying HDR content due to higher peak luminance
- Fewer degradation issues
- Provide additional power efficiency gains

Nits (cd/m2)

The Nit or “candela per square metre” (symbol: cd/m2) is the unit of luminance in the International System of Units (SI). The unit is based on the candela, the SI unit of luminous intensity, and the square metre, the SI unit of area. The nit (symbol: nt) is a non-SI name also used for this unit (1 nt = 1 cd/m2).
As a measure of light emitted per unit area, this unit is frequently used to specify the brightness of a display device. The sRGB spec for monitors targets 80 cd/m2.
Typically, monitors calibrated for SDR broadcast or studio color grading should have a brightness of 100 cd/m2.
HDR displays range from 450 to above 1600 cd/m2

OETF

Acronym of Opto-Electronic Transfer Function. Transfer function for converting brightness to electric signal on the camera side.

In post production, content is graded on a display that has a specific EOTF, tipically one that approximately reverses the camera’s OETF. (See also: EOTF and OOTF)

OLED

Stands for Organic Light-Emitting Diode. OLED TVs don't have a backlight in the traditional sense. Each individual pixel receives its own drive current and therefore can be individually controlled. OLEDs enable a TV to have a better contrast ratio as individual pixels can be switched off to obtain absolute black even while an adjacent pixel is at maximum brightness. This increases clarity whether you're standing far away or right next to it.
OLED displays can be subject to panel degradation and burn-in issues.

OOTF

Stands for Optical-to-Optical Transfer Function and constitutes "System Gamma". It is a mathematical function that maps scene luminance as seen by a camera to displayed luminance as produced by a monitor.

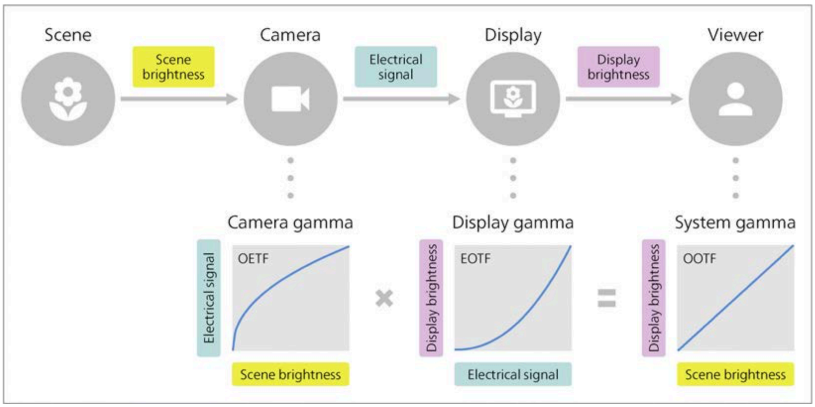

Credits: Leader Electronics Corporation Technical Information Series Vol. 02

Peak Code Value

Maximum digital code value that can be passed through a system component without clipping.

Peak Display Luminance

Highest luminance that a display can produce, in Nits.

PQ

Stands for Perceptual Quantizer and it is an EOTF (Electro-Optical Transfer Function).

It is based on the Barten curve (see page 9-10) and standardized in 2014 by the SMPTE as ST2084.

Perceptual quantization is an efficient way to encode High Dynamic Range (HDR) luminance's. Each consecutive pair of code values differ by just less than a perceivable step across the entire dynamic range, providing very efficient use of code values.

However, this EOTF does not offer backward compatibility for legacy displays, as PQ encoded signals can only be decoded by new HDR capable devices.

The PQ is designed for 10 and 12 bit content, and per the SMPTE ST 2084 standard, is not recommended for real-time broadcast.

Quantum Dot (QD) Displays

Quantum dot display is a display device that uses quantum dots (QD), semiconductor nanocrystals which can produce pure monochromatic[a] red, green, and blue light.
Quantum dots are excited by the blue light from the display panel to emit pure basic colors, which reduces light losses and color crosstalk in color filters, improving display brightness and color gamut. Light travels through QD layer film and traditional RGB filters made from color pigments, or through QD filters with red/green QD color converters and blue passthrough. The QD color filter technology is primarily used in LED-backlit LCDs but most recent version of Samsung's QD-OLED technology also uses Quantum Dot, pushing OLED picture quality (especially with regard to colour gamut coverage) even further.
Quantum dot displays are capable of displaying wider color gamuts, with some devices approaching full coverage of the BT.2020 color gamut.

Rec. 2020 or BT.2020

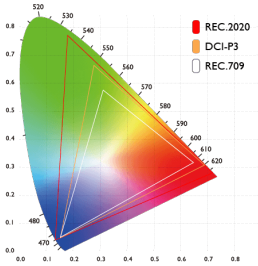

Credit: BenQ
Color Gamut: Understanding Rec.709, DCI-P3, and Rec.2020

ITU-R Recommendation BT.2020, more commonly known by the abbreviations Rec. 2020 or BT.2020, defines various aspects of ultra-high-definition television (UHDTV) with standard dynamic range (SDR) and wide color gamut (WCG), including picture resolutions, frame rates with progressive scan, bit depths, color primaries, RGB and luma-chroma color representations, chroma subsamplings, and an opto-electronic transfer function.
It defines a bit depth of either 10-bits per sample or 12-bits per sample.

The first version of Rec. 2020 was posted on the International Telecommunication Union (ITU) website on August 23, 2012,

and two further editions have been published since then.

Rec. 2020 is extended for high-dynamic-range (HDR) by Rec. 2100, which uses the same color primaries as Rec. 2020.

Rec. 2020 defines two resolutions: 3840 × 2160 (4K) and 7680 × 4320 (8K). These resolutions have an aspect ratio of 16:9 and use square pixels.

Rec. 2020 specifies the following frame rates: 120p, 119.88p, 100p, 60p, 59.94p, 50p, 30p, 29.97p, 25p, 24p, 23.976p. Only progressive scan frame rates are allowed.

Rec.709 or BT.709 or sRGB

Rec. 709, also known as Rec.709, BT.709, and ITU 709, is a standard developed by ITU-R for image encoding and signal characteristics of high-definition television.
The most recent version is BT.709-6 released in 2015. BT.709-6 defines the Picture characteristics as having an aspect ratio of 16:9, 1080 active lines per picture, 1920 samples per line, and a square pixel aspect ratio.
The first version of the standard was approved by the CCIR as Rec.709 in 1990, with the stated goal of a worldwide HDTV standard.

SEI Message

Stands for Supplemental Enhancement Information Message. The HEVC (High Efficiency Video Coding) version 2 adds several SEI messages:

Color Remapping Information: mapping one color space to another.
Knee Function: hints for converting between dynamic ranges, particularly from HDR to SDR.
Mastering Display Color Volume: describes the color primaries and dynamic range of the display that was used to author the video.
Time Code, mainly for archival purposes.

ST 2084

The perceptual quantizer (PQ), published by SMPTE as SMPTE ST 2084, is a transfer function that allows for HDR display by replacing the gamma curve used in SDR.
It is capable of representing luminance level up to 10000 cd/m2 (nits) and down to 0.0001 nits. It has been developed by Dolby and standardized in 2014 by SMPTE and also in 2016 by ITU in Rec. 2100.
ITU specifies the use of PQ or HLG as transfer functions for HDR-TV. PQ is the basis of HDR video formats (such as Dolby Vision, HDR10 and HDR10+) and is also used for HDR still picture formats. PQ is not backward compatible with the BT.1886 EOTF (i.e. the gamma curve of SDR), while HLG is compatible.
PQ is a non-linear transfer function based on the human visual perception of banding and is able to produce no visible banding in 12 bits. (*See Barten Curve, Page 9)*
A power function (used as EOTFs in standard dynamic range applications) extended to 10000 cd/m2 would have required 15 bits.

ST 2086

SMPTE ST2086 specifies the metadata items to specify the color volume (the color primaries, white point, and luminance range) of the display that was used in mastering video content. The metadata is specified as a set of values independent of any specific digital representation.

SDR

Standard-dynamic-range (SDR) video is a video technology which represents light intensity based on the brightness, contrast and color characteristics and limitations of a cathode ray tube (CRT) display. SDR video is able to represent a video or picture's colors with a maximum luminance around 100 cd/m2, a black level around 0.1 cd/m2 and Rec.709 / sRGB color gamut. It uses the gamma curve as its electro-optical transfer function.
The term "standard-dynamic-range video" was adopted to distinguish SDR video from high-dynamic-range video, a new technology that was developed in the 2010s to overcome SDR's limits

SBTM (HDMI)

Source-Based Tone Mapping or SBTM is a feature added to HDMI specification 2.1, revision a. It allows a TV and a source to exchange information on display performance, in such a way as to allow the source to carry out an optimized tone mapping of HDR content for the specific screen to which it is connected. It should be noted that this is not a new HDR format and will not involve any user activity.
Basically, with SBTM, tone mapping will no longer be done from

the display but from the source. Let's take for example a TV with a peak brightness of 800 cd/m². The screen will be able to communicate this data to the source, which at this point will adjust the output HDR signal in such a way as to allow an optimized reproduction.
(Source: DDay.it*)*

Tone Mapping

Is a technique used in image processing and computer graphics to map one set of colors to another to approximate the appearance of High Dynamic Range (HDR) images in a medium that has a more limited dynamic
range. Print-outs, CRT or Standard Dynamic Range (SDR) monitors, and projectors all have a limited dynamic range that is inadequate to reproduce the full range of light intensities present in HDR images. Tone mapping addresses the problem of strong contrast reduction from the recorded range to the displayable range while preserving the image details and color appearance, which are important to appreciate the original scene content and preserve creative intent. This tone mapping process is carried out using tone mapping operators, typically “S” shaped curves to roll off highlight and shadow detail. See also Inverse Tone Mapping (ITM) (p. 28).

Transfer Functions

The term "transfer function" in connection with HDR technology refers to the mathematical function or mapping that describes how the input values (brightness or luminance levels) of an image or video are transformed to the corresponding output values.

The transfer function in HDR can be represented by different mathematical models, such as the Perceptual Quantizer (PQ) curve or the Hybrid Log-Gamma (HLG) curve, which are commonly used in HDR standards like HDR10 and Dolby Vision.

A transfer function therefore describes the transformation of luminance values from the input to the output, enabling the capture, encoding, and display of high dynamic range content. Please see also Opto-Electrical Transfer Functions (OETF), used by cameras to convert visible light into an electric signal and their inverse, the Electro-Optical Transfer Functions (EOTF) used by displays to convert the electric signal into visible light patterns that we recognize as images. Combination of OETF and EOTF is referred to as the end-to-end Optical-Optical Transfer Function (OOTF).

Ultra HD

Stands for Ultra High Definition (also known as Super Hi-Vision, Ultra HDTV), as defined by the Consumer Technology Association (CTA), describes any display or content with an aspect ratio of at least 16:9 (1.77:1) and a resolution at least four times higher than "Full-HD" 1080p. 4K Ultra HD (2160p) and 8K Ultra HD (4320p) are two digital video formats proposed by NHK Science & Technology Research Laboratories and standardized by the International Telecommunication Union (ITU).

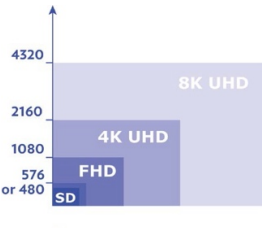

"4K" televisions have a resolution of 3,840 pixels wide by 2,160 pixels high (aka 2160p), while "8K" displays have a resolution of 7,680 pixels wide by 4,320 pixels high (4320p). "4K" panels feature four times the resolution of 1080p Full-HD displays.

These two formats utilize the 16:9 (1.77:1) aspect ratio, just like 720p and 1080p televisions.

Upscaling / Upconverting

In computer graphics and digital imaging, image scaling refers to the resizing of a digital image. In video technology, the magnification of digital material is known as upscaling or resolution enhancement.
When scaling a vector graphic image, the graphic primitives that make up the image can be scaled using geometric transformations, with no loss of image quality. When scaling a raster graphics image, a new image with a higher or lower number of pixels must be generated. In the case of decreasing the pixel number (scaling down) this usually results in a visible quality loss. From the standpoint of digital signal processing, the scaling of raster graphics is a two-dimensional example of sample-rate conversion, the conversion of a discrete signal from a sampling rate (in this case the local sampling rate) to another.

VP-9

VP9 is an open and royalty-free video coding format developed by Google.
VP9 is the successor to VP8 and competes mainly with MPEG's High Efficiency Video Coding (HEVC/H.265). At first, VP9 was mainly used on Google's video platform YouTube.The emergence of the Alliance for Open Media, and its support for the ongoing development of the successor AV1, of which Google is a part, led to growing interest in the format.
In contrast to HEVC, VP9 support is common among modern web browsers. Android has supported VP9 since version 4.4

KitKat, while iOS/iPadOS added support for VP9 in iOS/iPadOS 14.
Parts of the format are covered by patents held by Google. The company grants free usage of its own related patents based on reciprocity, i.e. as long as the user does not engage in patent litigations.

VVC (H.266)

Versatile Video Coding (VVC), also known as H.266, ISO/IEC 23090-3, and MPEG-I Part 3, is a video compression standard finalized on 6 July 2020, by the Joint Video Experts Team (JVET). It is the successor to High Efficiency Video Coding (HEVC, also known as ITU-T H.265 and MPEG-H Part 2). It was developed with two primary goals – improved compression performance and support for a very broad range of applications.

VVC is not royalty free. Four companies were initially vying to be patent pool administrators for VVC, in a situation similar to the previous AVC and HEVC codecs. Two companies were later reported to be forming pools: Access Advance and MPEG LA.

Wide Color Gamut (WCG)

Stands for Wide Color Gamut. Its includes colors significantly more saturated than those that can be represented using Recommendation ITU-R BT.709, such as the color space defined in Rec. 2020.

White Point

A white point (often referred to as reference white or target white in technical documents) is a set of chromaticity coordinates that serve to define the color “white” in image capture, encoding, or reproduction. Depending on the application, different definitions of “white” are needed to give acceptable results. For example, photographs taken indoors may be lit by incandescent lights, which are relatively orange compared to daylight. Therefore most professional cameras have different settings for shooting under incandescent lighting vs. daylight. Likewise, images that are meant to be viewed on a display with a “D65” white point will appear incorrect on a display with a different white point.

CIE standard illuminant D65 is frequently used to define the white point for video displays.

D55 was the standard white point for film projection. Both the “DCI” white point and D60 are common for many digital cinema motion motion pictures.

XML

Extensible Markup Language (XML) is a computer language used to describe data that defines a set of rules for encoding documents in a format that is readable by both humans and machines. The XML can be delivered as a separate “side car” file or embedded in an IMF Package.
In case of Dolby Vision, multiple levels of metadata can be present. A full description of Dolby Vision Metadata can be found here:

https://professionalsupport.dolby.com/s/article/Dolby-Vision-Metadata-Levels?language=en_US

Reference of International Standards and Specifications

ST 2084:2014 - SMPTE Standard - High Dynamic Range Electro- Optical Transfer Function of Mastering Reference Displays
ST 2086:2018 - SMPTE Standard - Mastering Display Color Volume Metadata Supporting High Luminance and Wide Color Gamut Images
ITU-R BT.2100-2, “Image parameter values for high dynamic range television for use in production and international programme exchange,” July 2018.
ITU-R BT.2390-7, “High dynamic range television for production and international programme exchange,” July 2019.
ITU-R BT.2408-1, “Operational practices in HDR television production,” April 2018.
ITU-R BT.814-4, “Specifications of PLUGE test signals and alignment procedures for setting of brightness and contrast of displays,” July 2018.
ITU-R BT.2111-0, “Specification of colour bar test pattern for high dynamic range television systems,” December 2017.
ITU-R BT.2020-2, “Parameter values for ultra-high definition television systems for production and international programme exchange,” October 2015.
ITU-R BT .2246-6, “The present state of ultra-high definition television, ” March 2017.
CTA-861G, “A DTV Profile for Uncompressed High Speed Digital Interfaces, ” November 2016.
HDMI Specification Ver.2.0b, “High-Definition Multimedia Interface. ”
AV1 Bitstream & Decoding Process Specification (GitHub: https://aomediacodec.github.io/av1-spec/av1-spec.pdf)
ISO/IEC 23094-2:2021 – General video coding — Part 2: Low complexity enhancement video coding
ITU- T H.265, “High efficiency video coding, ” February 2018.
ITU-T H.266 (04/2022) H.266 Versatile Video Coding

Sponsored by

Grass Valley EIZO telestream

The HDR Fact Book

Third Edition - March 2025

Editor and content supervision: Alessandro Capitani

For educational purposes only, not for sale

Video Progetti
A feeling for technology.

Video Progetti Srl - Via Cosimo Tornabuoni 21, 00166 Roma (RM) - +39 066129641

Via Rivoltana 35, 20096 Pioltello (MI) - +39 0289305809