Video and Colorspaces

Hans Verkuil
Cisco Systems Norway

What Is Color?

Spectral Power Distribution

- Power per area per wavelength: i.e. the intensity of the light of different wavelengths.

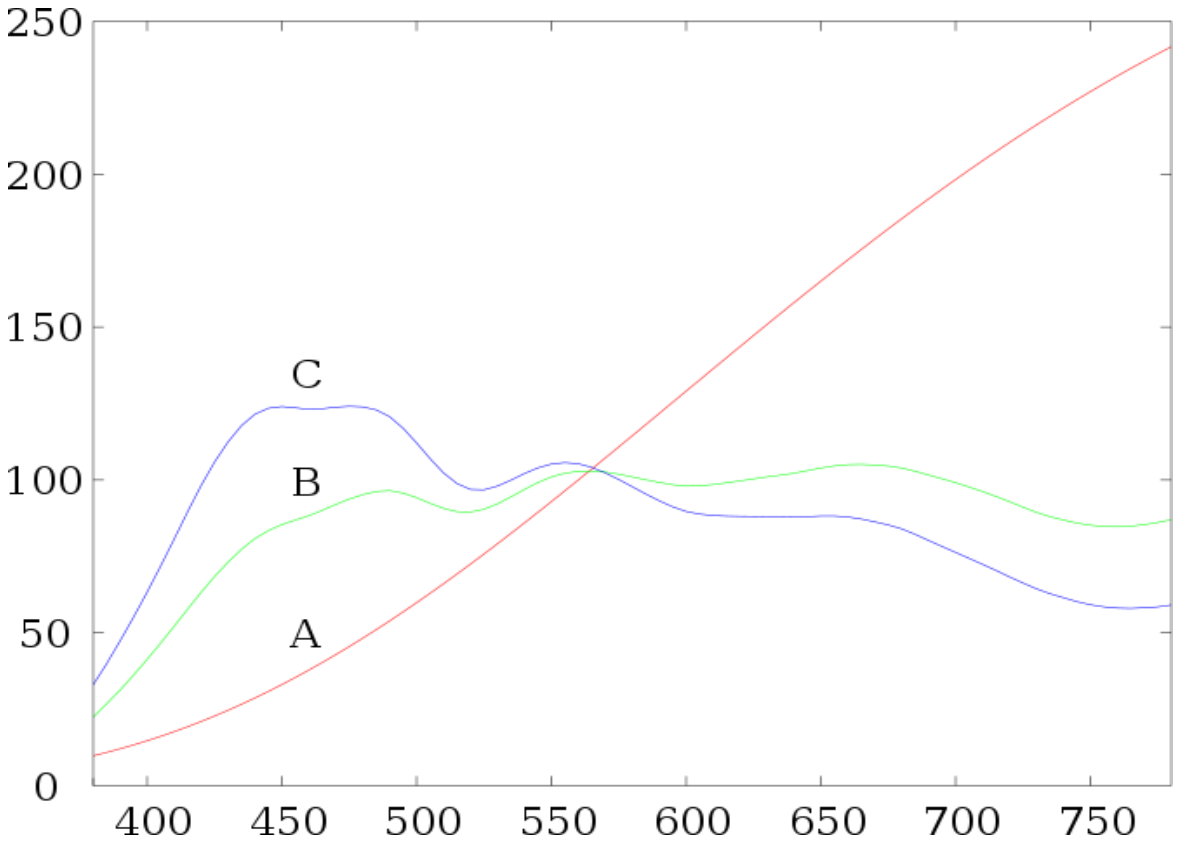

"CIE illuminants A,B,C" by Adoniscik - Own work. Licensed under Public domain via Wikimedia Commons - http://commons.wikimedia.org/wiki/File:CIE_illuminants_A,B,C.svg#mediaviewer/File:CIE_illuminants_A,B,C.svg

The Eye

- Three types of cones detect color: S, M and L cones with peak sensitivities around 440 nm, 545 nm and 565 nm. The visible spectrum is about 390 – 700 nm.

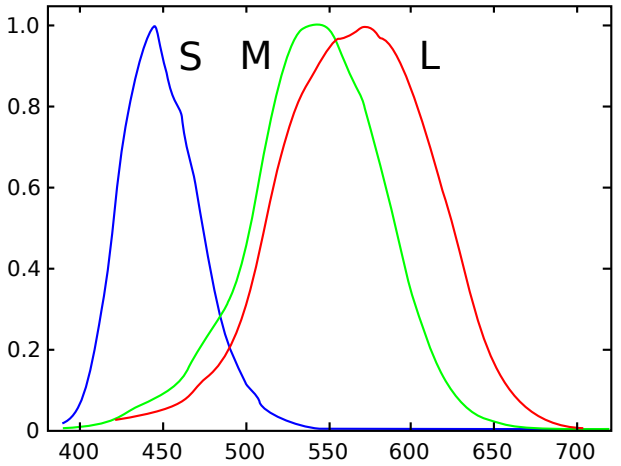

"Cones SMJ2 E" by Vanessaezekowitz at en.wikipedia / Later version uploaded by BenRG. - Based on Dicklyon's PNG version, itself based on data from Stockman, MacLeod & Johnson (1993) Journal of the Optical Society of America A, 10, 2491-2521d http://psy.ucsd.edu/~dmacleod/publications/61StockmanMacLeodJohnson1993.pdf (log E human cone response, via http://www.cvrl.org/database/text/cones/smj2.htm)Transferred from en.wikipedia to Commons by User:Richard001 using CommonsHelper.. Licensed under Creative Commons Attribution-Share Alike 3.0-2.5-2.0-1.0 via Wikimedia Commons - http://commons.wikimedia.org/wiki/File:Cones_SMJ2_E.svg#mediaviewer/File:Cones_SMJ2_E.svg

An Illusion

- Light from an object with a certain Spectral Power Distribution hits the cones in the eye and the SPD is reduced to three 'values', one from each cone type.
- These three values are interpreted as colors by the brain. Many different SPDs will all map to the same three values.
- To reproduce colors all you have to do is create an SPD that results in the same three values coming from the cones.
- Displays and photos do not recreate the original image, they just recreate the cone impulses to the brain.
- Three cones, so you need three colors around the cone peak sensitivities to recreate the illusion.
- The mantis shrimp is sensitive to 12 different colors, and so would need displays with 12 colors to recreate the illusion.

An Illusion

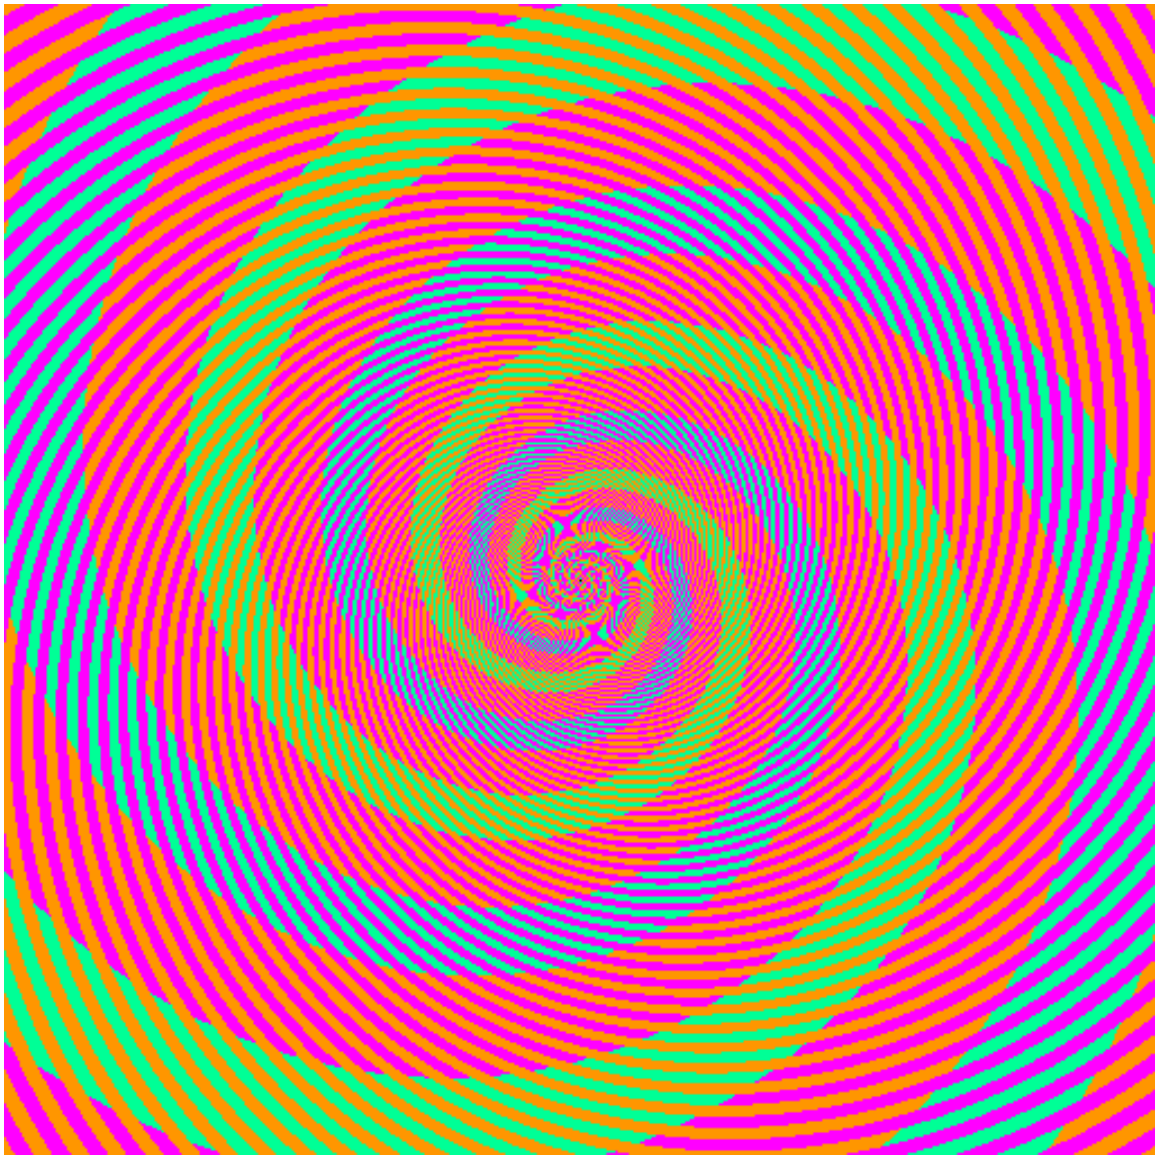

An Illusion

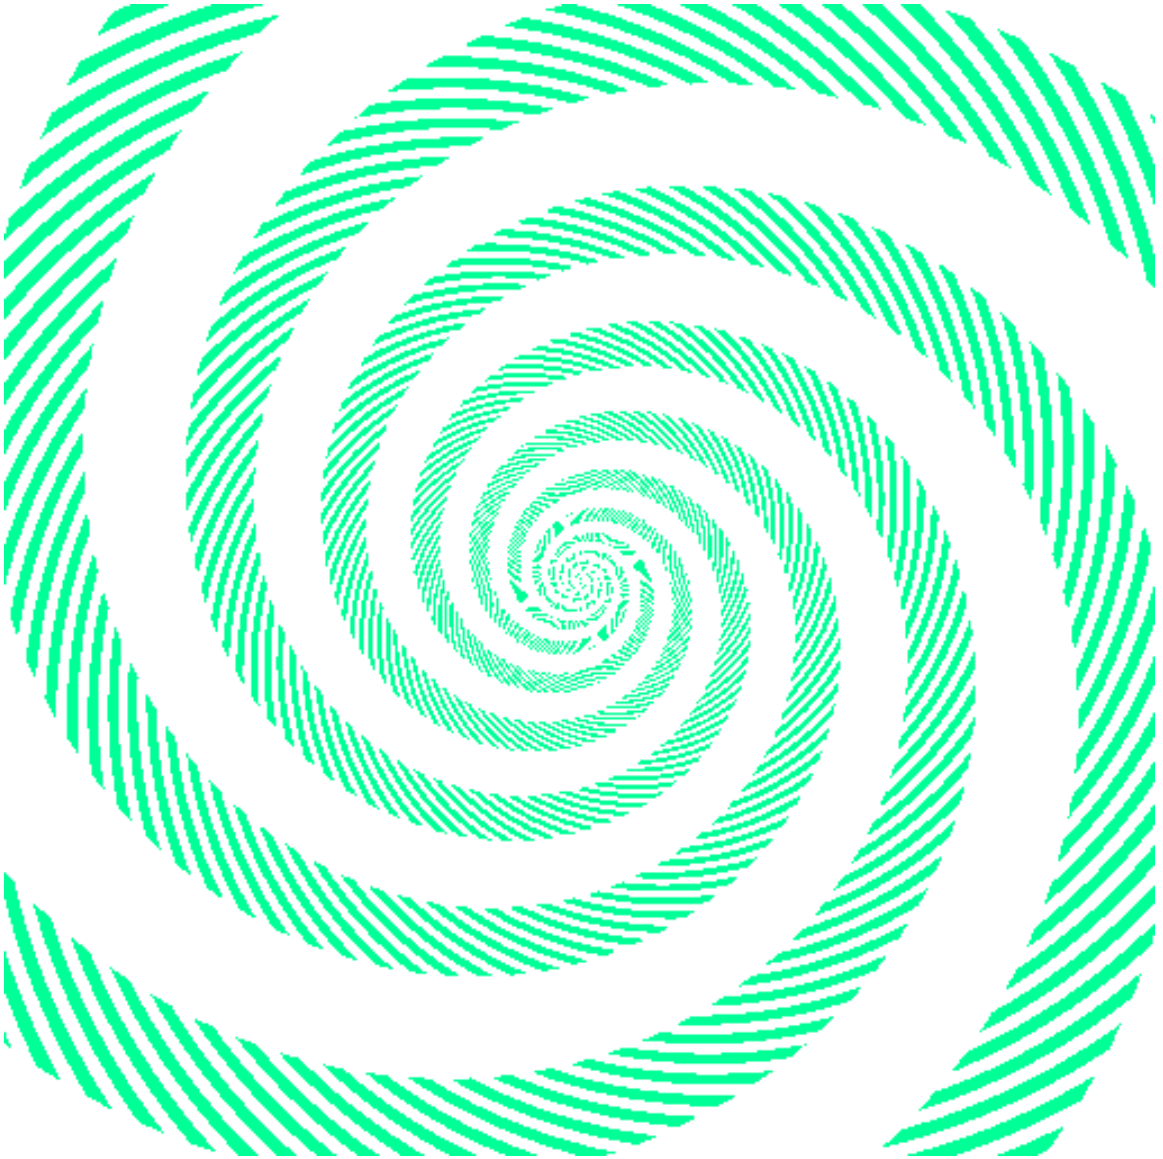

CIE XYZ

- Studies done in the 1920s resulted in data sets that determined how a specific wavelength can be reproduced using three primary light sources of 435.8, 546.1 and 700 nm. Also known as the CIE RGB colorspace.
- CIE: *Commission internationale de l'éclairage* or *International Commission on Illumination*.
- In practice the CIE XYZ colorspace is used, which is a linear transformation of RGB to an all positive set of color matching functions using imaginary primaries (i.e. light sources that are not physically possible).
- CIE XYZ is the foundation of all colorspaces, based from measurements taken from a pool of 17 people over 85 years ago!
- Note: Y describes luminance.

Colorspaces

- CIE xyz: $x = X / (X + Y + Z)$,
  $y = Y / (X + Y + Z)$,
  $z = Z / (X + Y + Z) = 1 - x - y$
- CIE Yxy: Y = luminance, xy = chromaticity.
- A colorspace defines which three primaries are used to recreate colors. These are usually defined as a (x, y) tuple. These tuples represent the color vectors (1, 0, 0), (0, 1, 0), (0, 0, 1).
- A colorspace also defines the white point (x, y) tuple, which is the color obtained by the maximum of all three primaries or the color vector (1, 1, 1). It effectively defines the relative power or energy of the primaries.
- In the past these were based on phosphors used in TVs.
- Any colorspace can be derived from the CIE XYZ colorspace through a 3x3 matrix multiplication.

Colorspaces

- All colorspaces can represent all colors (although with values outside the 0-1 range), but not after quantization, i.e. when values < 0 or > 1 are cut off. This defines the gamut or extent of a colorspace.
- Colorspaces are linear but early CRTs weren't and neither is human vision: double the luminance will not actually look twice as bright.
- Colorspaces define a transfer function to go from linear to a non-linear color representation. Denoted as R'G'B' where each component has gone through the transfer function. The transfer function is sometimes called 'gamma' which is incorrect.
- Y'CbCr is in turn derived from R'G'B' via a matrix multiplication. Colorspaces can define one or more possible matrices for this. Note that Y'CbCr is *not* a colorspace, it is just a different way of representing colors within a colorspace.
- Take care when using textures in openGL: openGL expects linear RGB by default, not R'G'B'!

Color Transformations

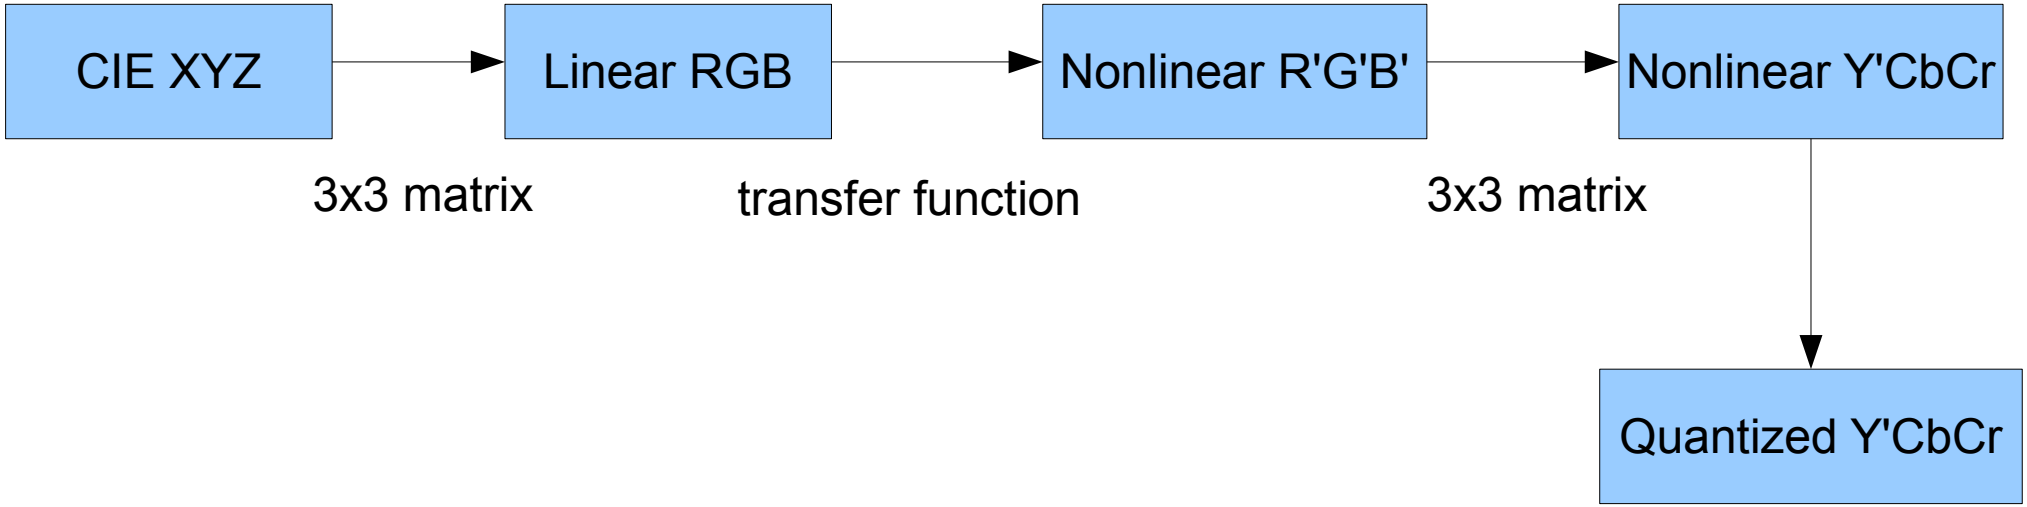

Standards

- Best known: sRGB. Used for computer graphics. Standard: IEC 61966-2-1:1999. Chromaticities identical to Rec. 709, but different transfer function. IEC 61966-2-1 Amendment 1 defines the sYCC Y'CbCr encoding based on BT.601.
- SMPTE 170M (aka SMPTE C): defines the SDTV colorspace. The transfer function is identical to Rec. 709, but the chromaticities are different. The Y'CbCr encoding is the same as in BT.601. Note: BT.601 only defines the Y'CbCr encoding, not a colorspace. BT.601 is sometimes – incorrectly – used as an alias for SMPTE 170M. An extended gamut Y'CbCr encoding (xvYCC 601) can optionally be used (IEC 61966-2-4).
- Rec. 709: HDTV colorspace. Standard: Rec. ITU-R BT.709-5. An extended gamut Y'CbCr encoding (xvYCC 709) can optionally be used (IEC 61966-2-4).
- AdobeRGB: the Y'CbCr encoding follows BT.601. Also called opRGB (IEC 61066-2-5), but with a slightly different gamma (2.2 instead of 2.19921875) as defined by the Adobe RGB 1998 standard.
- BT.2020: used for UHDTV with deep color (>= 10 bits per component). The transfer function is the same as Rec. 709. It defines two different Y'CbCr encodings.

Standards

- SMPTE 240M: obsolete, used by HDTV equipment from 1988-1998 when it was superseded by Rec. 709.
- NTSC 1953: the original NTSC colorspace. Superseded by SMPTE 170M.
- EBU Tech. 3213: the original PAL/SECAM colorspace. Superseded by SMPTE 170M.
- Studios are moving towards Rec. 709 chromaticities for SDTV instead of those defined by SMPTE 170M. Unfortunately, there usually is no way of knowing what colorspace is used for SDTV.

Limited and Full Range

- Typically 8-bit R', G' and B' values are quantized to the range [0-255]. This is full range quantization.
- Typically 8-bit Y' values are quantized to the range [16-235] and Cb and Cr values to the range [16-240]. This is limited range quantization.
- But limited range R'G'B' (values in the range [16-235]) and full range Y'CbCr variants exist as well. HDMI can signal both variants.

HDMI (CEA-861-F)

- If the sink cannot receive AVI InfoFrames or the sink cannot receive Y'CbCr then:
  - for IT (i.e. non-SDTV/HDTV) timings use full range sRGB.
  - for SDTV/HDTV timings use limited range sRGB. God knows why...
- else:
  - for IT (ITC=1) (or non-SDTV/HDTV) timings use full range sRGB.
  - for SDTV timings use limited range Y'CbCr SMPTE 170M.
  - for HDTV timings use limited range Y'CbCr Rec. 709.
- The AVI InfoFrame supports colorimetry information that can signal:
  - SMPTE 170M, Rec 709, xvYCC using BT.601 Y'CbCr encoding, xvYCC using Rec. 709 Y'CbCr encoding, AdobeRGB using BT.601 Y'CbCr encoding, AdobeRGB, BT.2020 Y'cCbcCrc (Constant Luminance), BT.2020 Y'CbCr, BT.2020 RGB.
  - Limited vs Full range RGB, Limited vs Full range Y'CbCr.

Guidelines

- SDTV: SMPTE 170M.
- HDTV: Rec. 709.
- Graphics: sRGB.

Problems

- Highly confusing colorspace names.
- Inconsistent conversion matrices: often buggy, never the same.
- Applications often ignore colorspace information.
- Incorrect handling of limited/full range quantization.
- Hardware does not take different transfer functions into account. E.g. the adv7604 driver produces the wrong results when converting from Rec. 709 Y'CbCr to sRGB R'G'B'.

V4L2 API Enhancements in 3.19

- Added two new colorspaces: ADOBERGB and BT2020.
- Added a new ycbcr_enc field to v4l2_pix_format(_mplane) to describe Y'CbCr encoding variants:
  - V4L2_YCBCR_ENC_DEFAULT: use the colorspace's default encoding
  - V4L2_YCBCR_ENC_601: BT.601 encoding
  - V4L2_YCBCR_ENC_709: Rec. 709 encoding
  - V4L2_YCBCR_ENC_XV601: xvYCC BT.601 extended gamut encoding
  - V4L2_YCBCR_ENC_XV709: xvYCC Rec. 709 extended gamut encoding
  - V4L2_YCBCR_ENC_SYCC: sYCC encoding (Y'CbCr variant of sRGB)
  - V4L2_YCBCR_ENC_BT2020: BT.2020 Non-Constant Luminance encoding
  - V4L2_YCBCR_ENC_BT2020_CONST_LUM: BT.2020 Constant Luminance encoding
  - V4L2_YCBCR_ENC_SMPTE240M: Obsolete SMPTE240M Y'CbCr encoding
- Added a new quantization field to v4l2_pix_format(_mplane) to describe difference quantization methods:
  - V4L2_QUANTIZATION_DEFAULT: full range for R'G'B', xvYCC and sYCC, limited range otherwise.
  - V4L2_QUANTIZATION_FULL_RANGE: full range.
  - V4L2_QUANTIZATION_LIM_RANGE: limited range.

V4L2 API Enhancements

- For DisplayPort 1.3 we miss: scRGB (floating point), RGB wide gamut fixed point (XR8, XR10, XR12) and DCI-P3 (Cinema projector colorspace).
  - It is unclear whether scRGB is linear or non-linear (probably linear).
  - I have no idea what the fixed point RGB colorspace refers to.
  - DCI-P3 is standardized in SMPTE-431-2 and could be added as V4L2_COLORSPACE_DCI_P3.
  - I cannot test these exotic colorspaces on DisplayPort.

Resources

- Color Imaging – Fundamentals and Applications, Erik Reinhard et. al.
- Digital Video and HDTV – Algorithms and Interfaces, Charles Poynton.
- http://www.brucelindbloom.com
- http://hverkuil.home.xs4all.nl/spec/media.html#colorspaces

Questions?