High Dynamic Range Imaging

Is there a role for ICC?

W Craig Revie

NOTE: the ideas expressed in this presentation do not yet have industry consensus. While I have taken care to check the details presented, some of these may be incorrect and all are subject to change.

Starting assumptions

- HDR displays will become increasingly popular for computers and mobile devices and at least some of these displays will have a high luminance (500 cd/m2 or more)
- SDR content must be restricted to use a lower luminance white for these displays to avoid user discomfort (and possibly screen fading) and many displays further restrict the display of HDR by applying power limits
- Some HDR content will be prepared for display on a range of HDR displays and in many cases the content provider will have no control over how the content will be displayed (a ‘broadcast’ model)
- Users of HDR displays will wish to be able to mix HDR and SDR content, for example on a web browser, a PDF reader and other document editors such as MS Word and there is currently not a good solution that will support this
- The broadcast model requires a common exchange space for the content and a mapping for each display (or display category)

My initial conclusion

- Some of this functionality can be provided by ICC v4 ICC Profiles with a minor extension (Adobe’s HDR extension/ bug fix) by adopting the HLG model (this may compromise the quality of some PQ-based HDR content)
- iccMAX may be able to provide a more elegant and complete solution
- High quality display of HDR content may require a more general model which provides support for dynamic metadata

Example use-cases

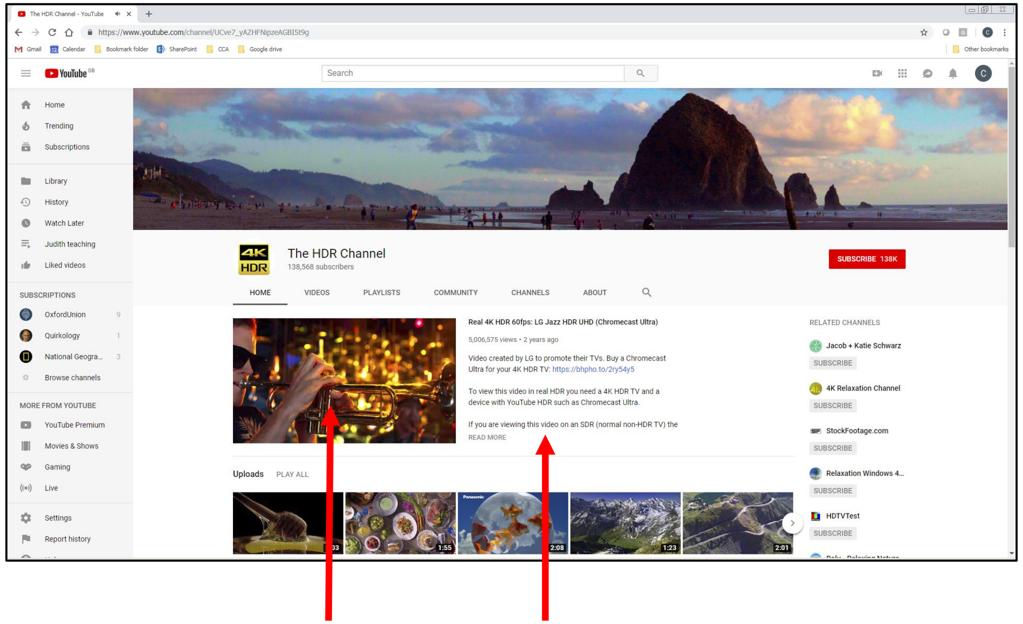

HDR SDR

Web browser

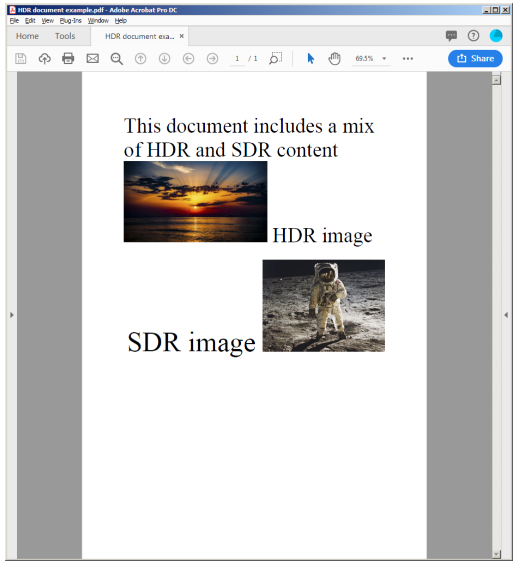

Document creation and display

Printing for backlit display (and printing in general)

References

- ***Recommendation ITU-R BT.2100-2 (07/2018)***, Image parameter values for high dynamic range television for use in production and international programme exchange.
  See https://www.itu.int/dms_pubrec/itu-r/rec/bt/R-REC-BT.2100-2-201807-I!!PDF-E.pdf
- ***Report ITU-R BT.2408-1 (04/2018)***, Operational practices in HDR television production.
  See https://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-BT.2408-1-2018-PDF-E.pdf
- ***Report ITU-R BT.2390-4 (04/2018)***, High dynamic range television for production and international programme exchange.
  See https://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-BT.2390-4-2018-PDF-E.pdf

Two main approaches (see BT.2100)

Hybrid Log Gamma (HLG)

- Promoted by BBC and NHK
- Designed for live broadcast
- No metadata used
- Fits well with the ICC model but requires a small extension and rethinking of the concept of Relative Colorimetric rendering

Perceptual Quantisation (PQ)

- Promoted by Dolby and many others
- Designed for motion picture production
- Variants include Dolby Vision, ***HDR10***, HDR10+ all of which use metadata
  - Maximum Frame–Average Light Level (MaxFALL)
  - Maximum Content Light Level (MaxCLL)
  - Colour gamut
- Does not fit with the ICC model as communication is required between source and destination profiles

HLG HDR System

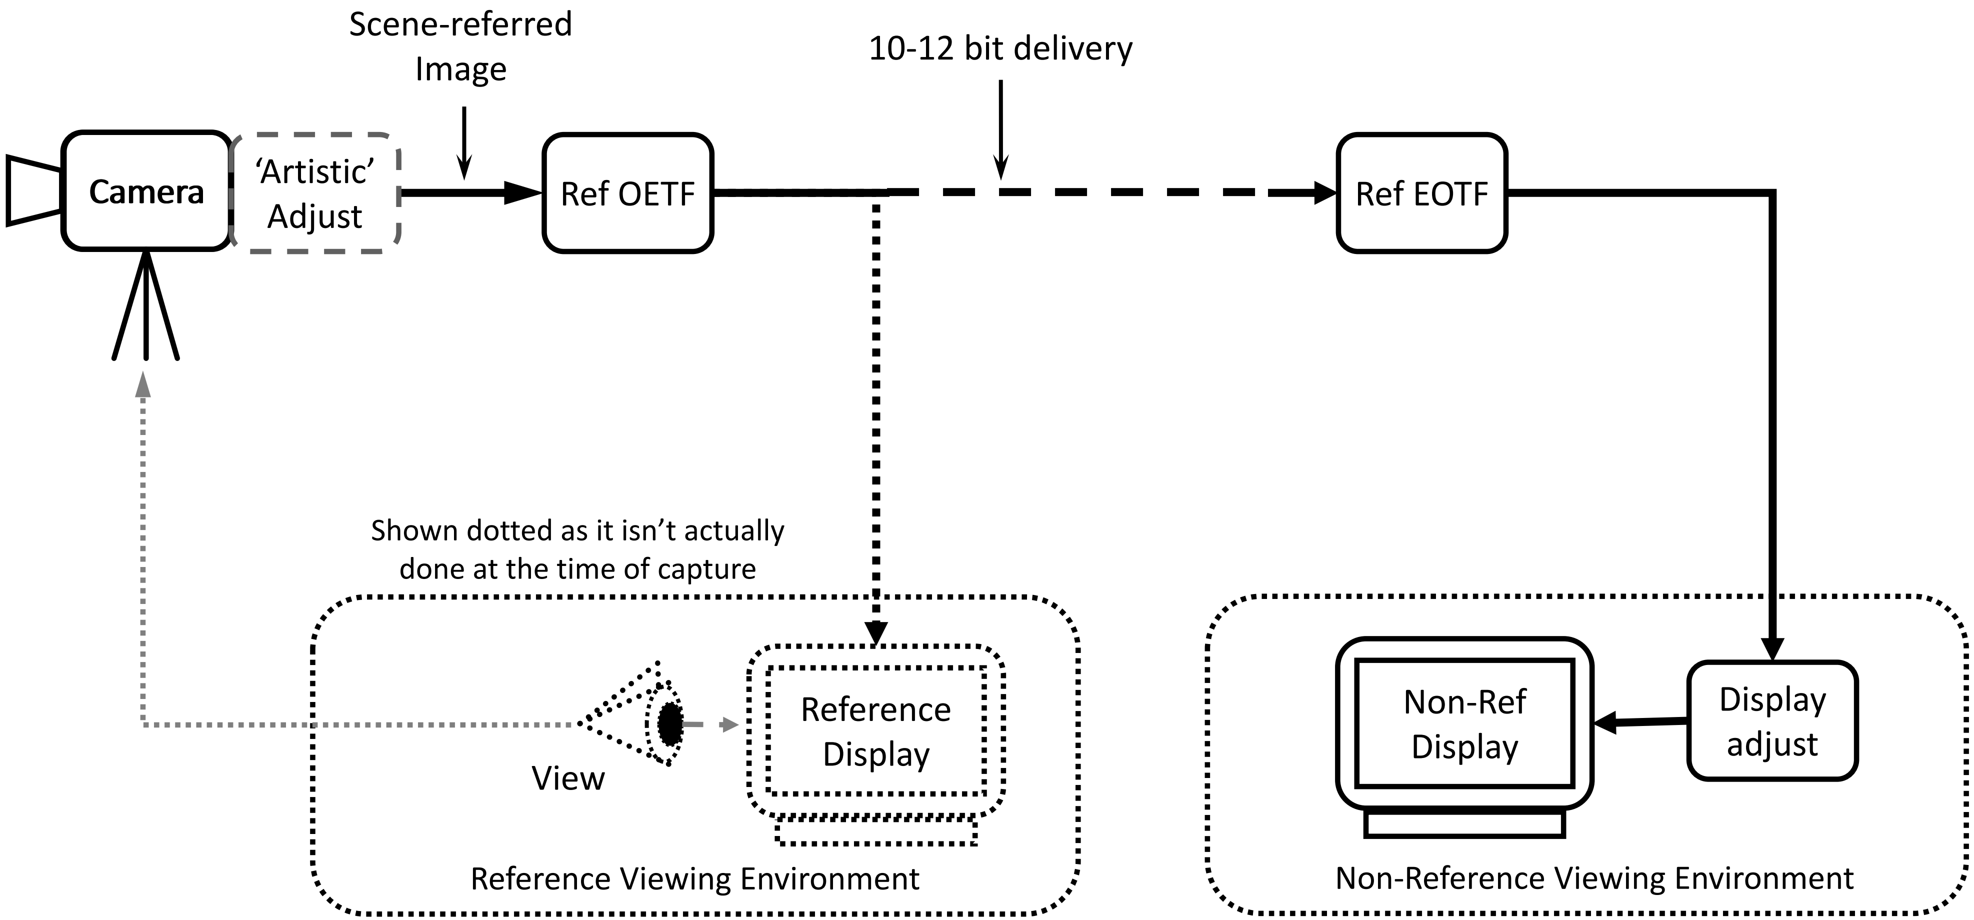

Based on ITU-R BT.2390-4 FIGURE 10

HLG reference broadcast model based on Report ITU-R BT.2390-5 (10/2018) and Recommendation ITU-R BT.2100-2 (07/2018)

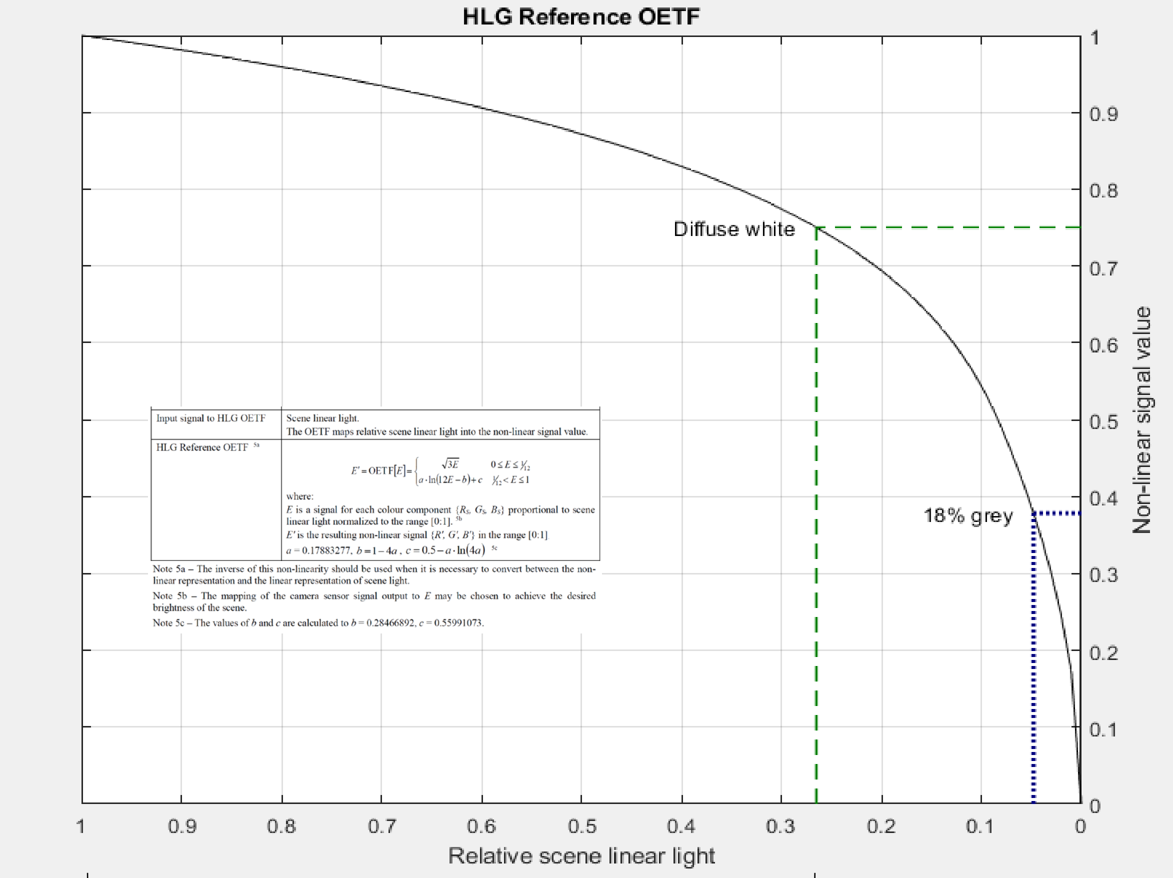

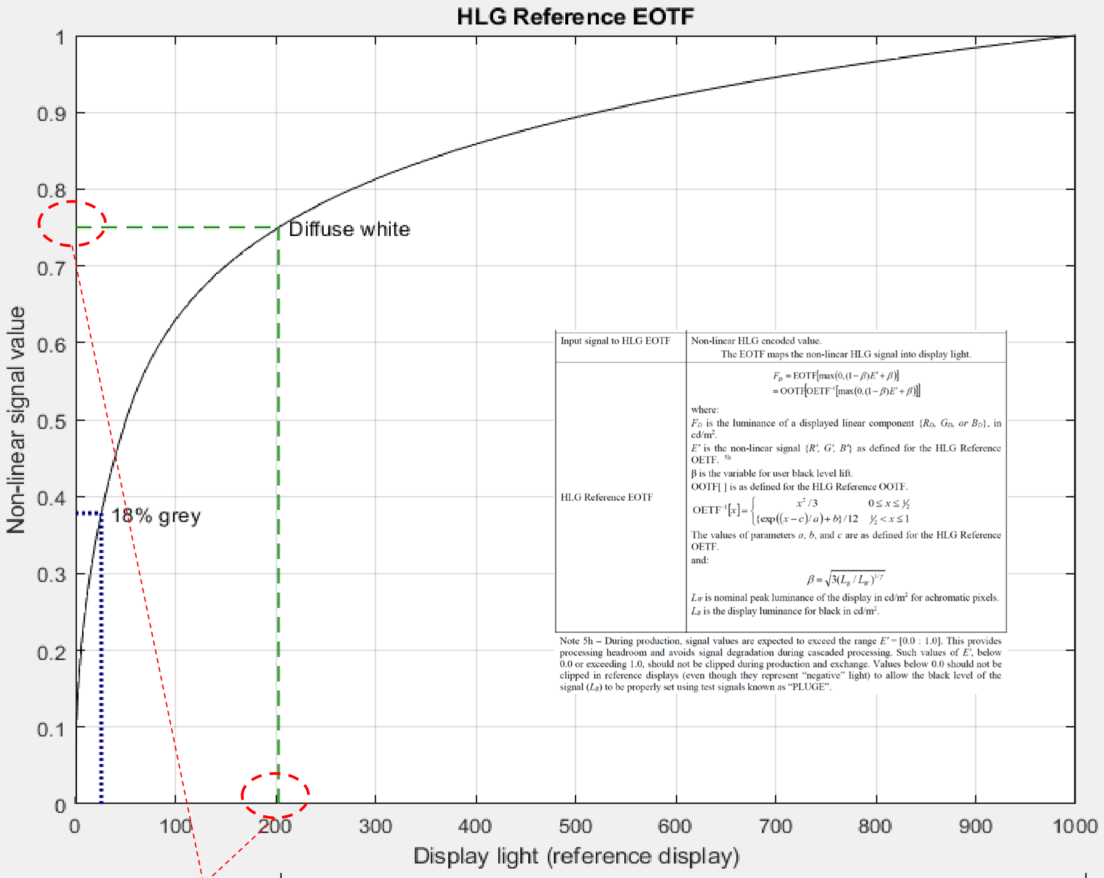

Peak white is
3.8 x diffuse white

Peak white is
5 x diffuse white

Nominal signal levels for PQ and HLG production

Reflectance Object or Reference (Luminance Factor, %)[3]	Nominal Luminance, cd/m² (PQ & 1000 cd/m² HLG)	Nominal Signal Level	
		%PQ	%HLG
Grey Card (18%)	26	38	38
Greyscale Chart Max (83%)	162	56	71
Greyscale Chart Max (90%)	179	57	73
Reference Level: HDR Reference White (100%) also diffuse white and Graphics White	203	58	75

TABLE 1 from Report ITU-R BT.2408-1

HLG with ICC

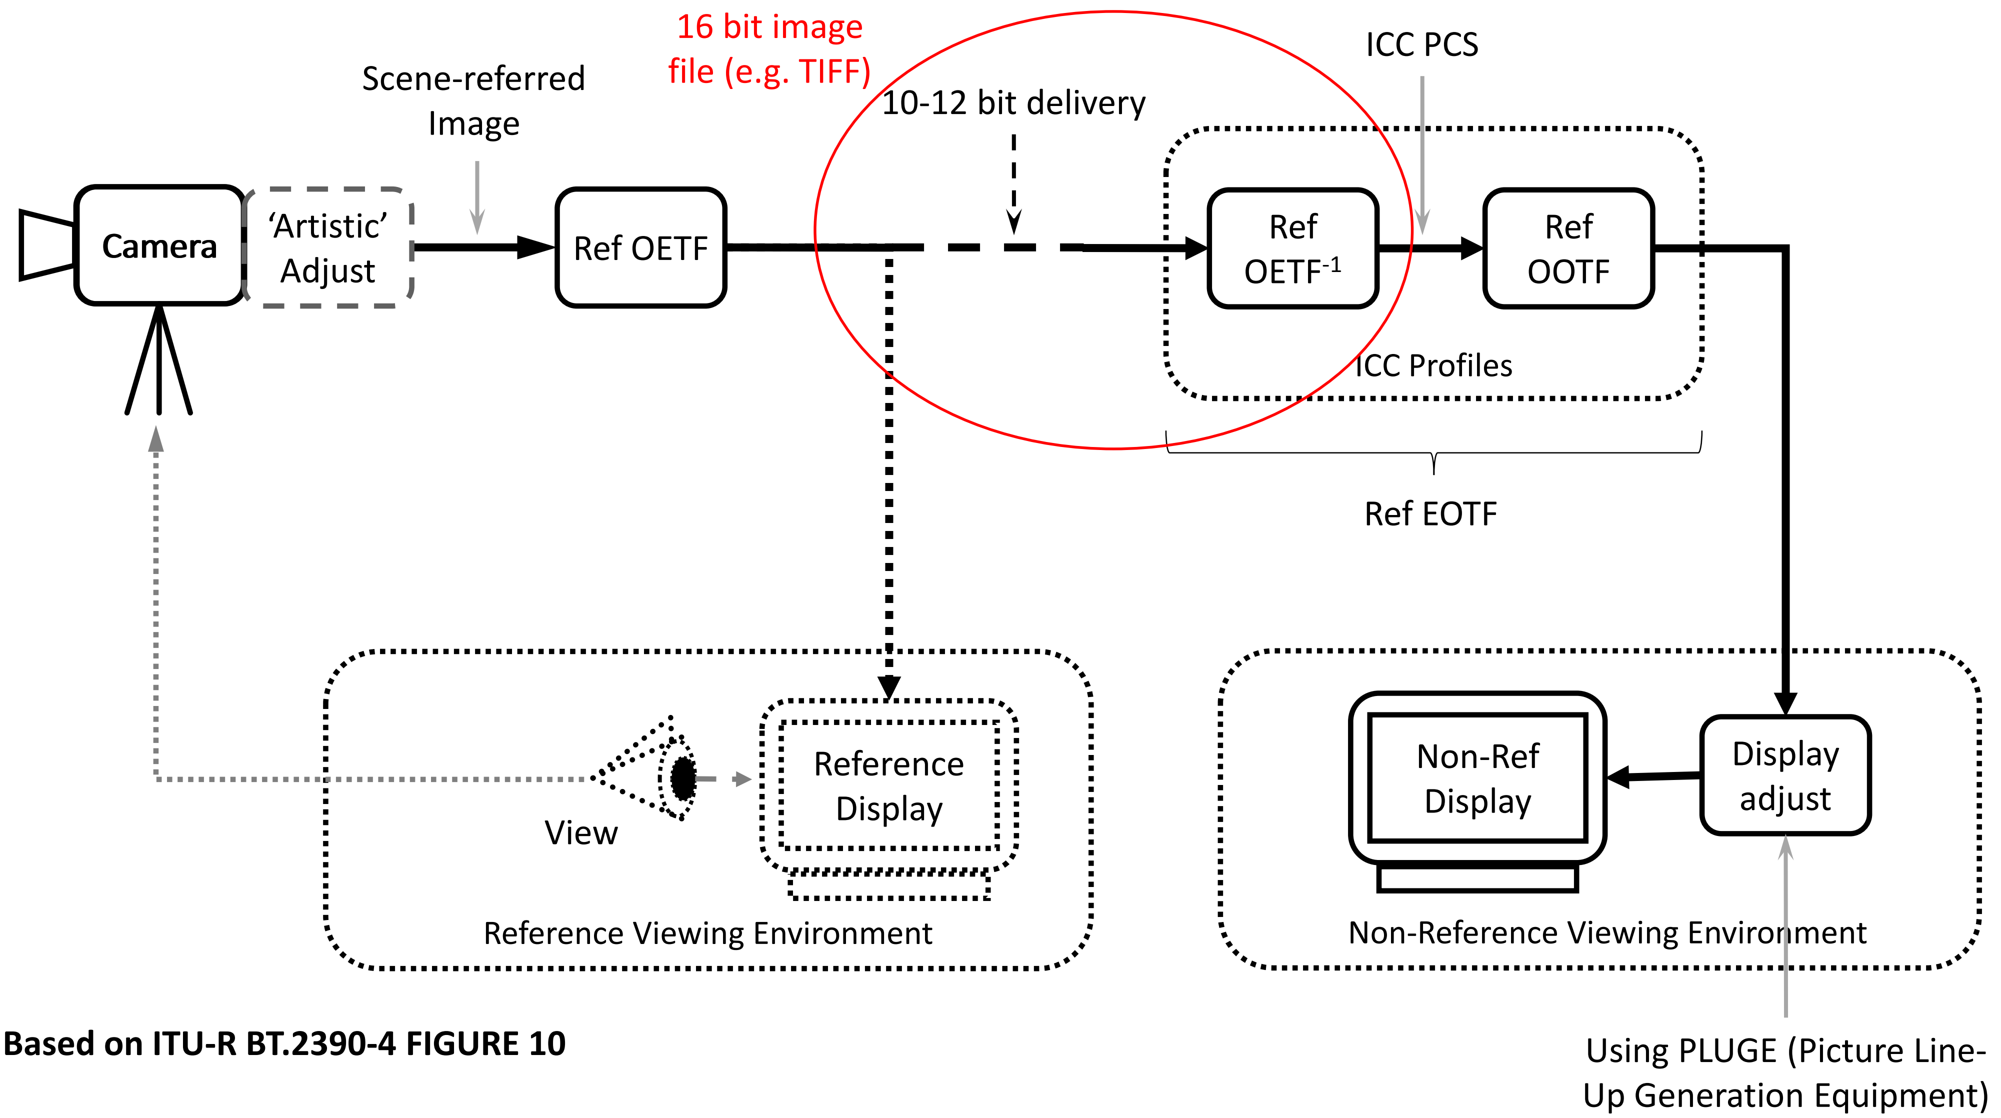

Based on ITU-R BT.2390-4 FIGURE 10

ICC colour management for HLG

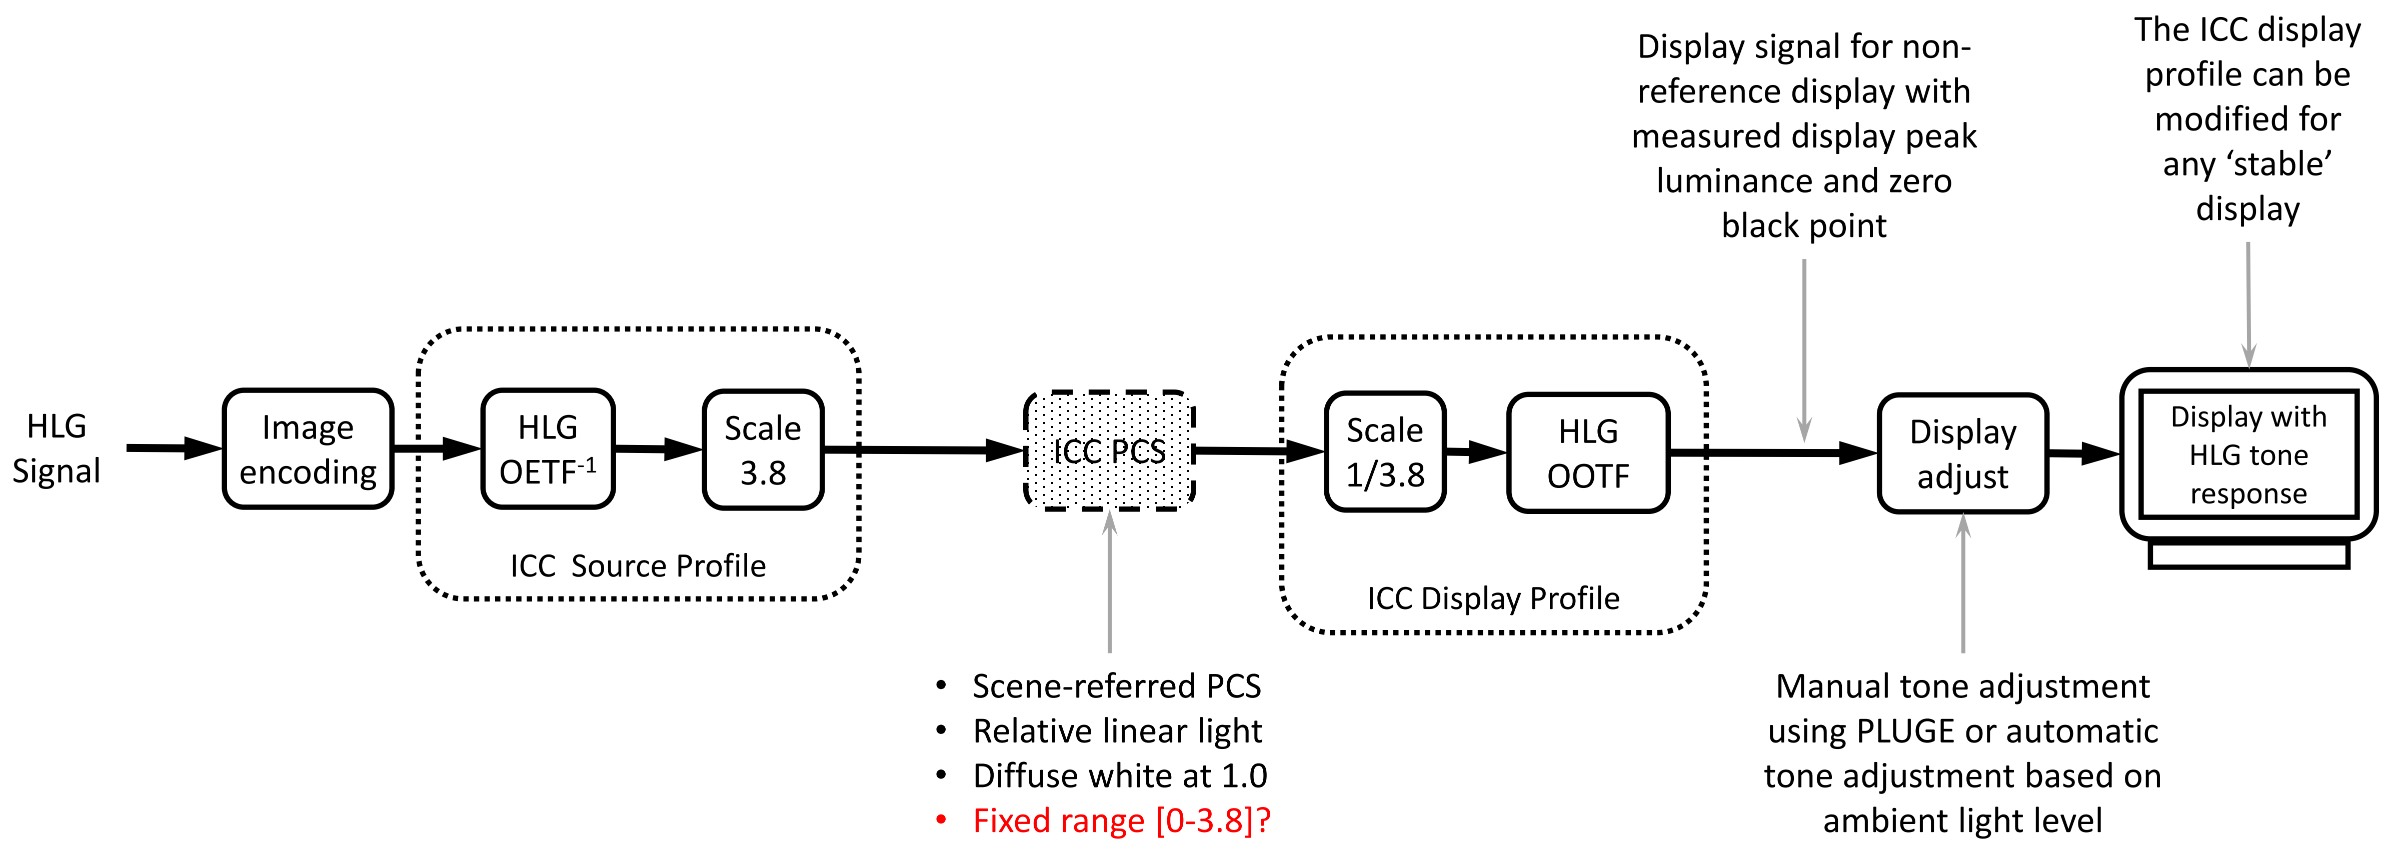

Possible ICC HDR support for HLG based on Report ITU-R BT.2390-4 (04/2018) and Recommendation ITU-R BT.2100-2 (07/2018)

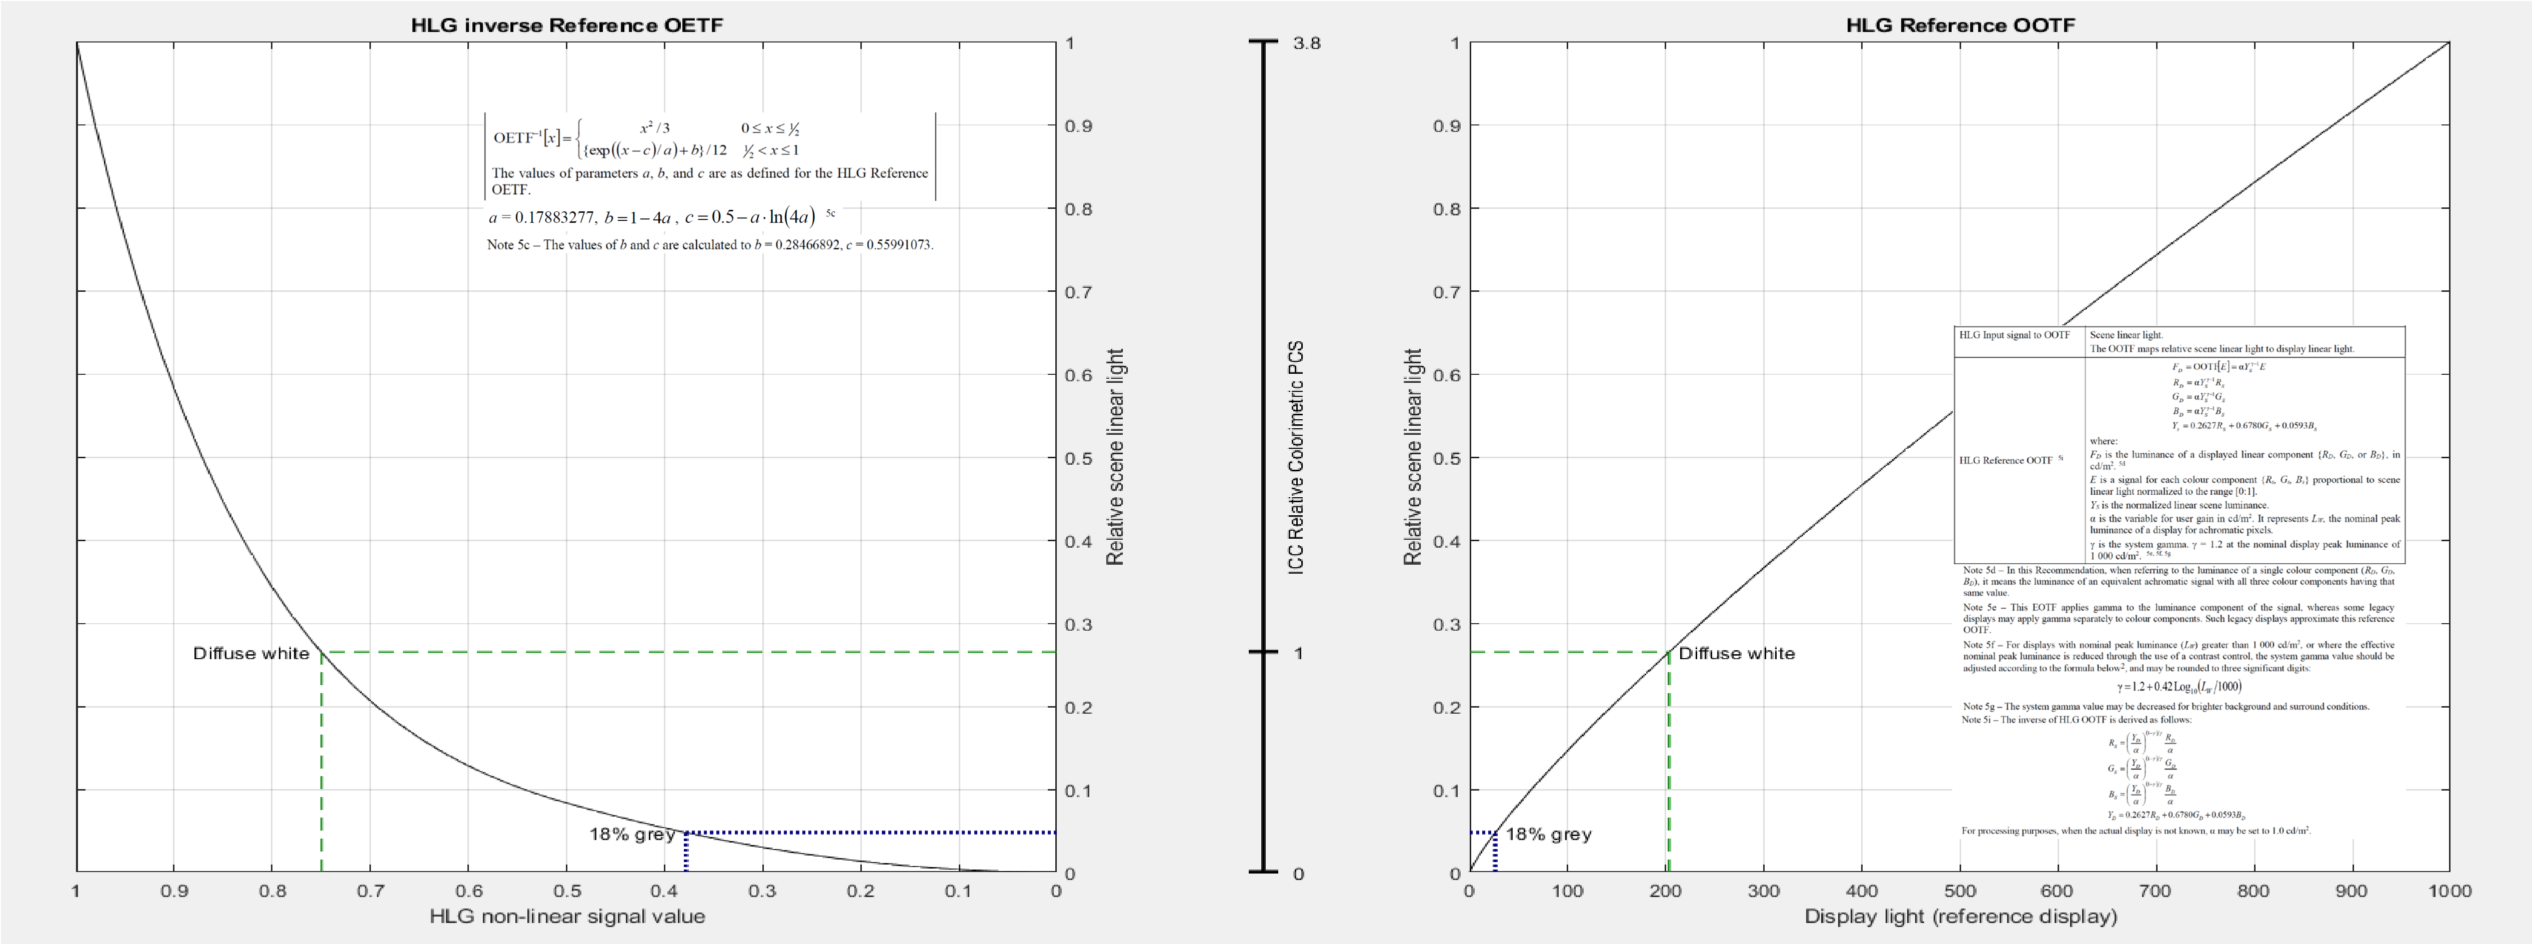

Lars Borg (Adobe) created an HLG Scene ICC Profile that incorporates this transform (Rec2100HLG203Scene.icc)

- An ICC profile that incorporates this transform will be required in order to convert images for a reference display
- A variant of this transform will be required where the display peak white or black luminance differs from that of the reference display

PQ HDR System

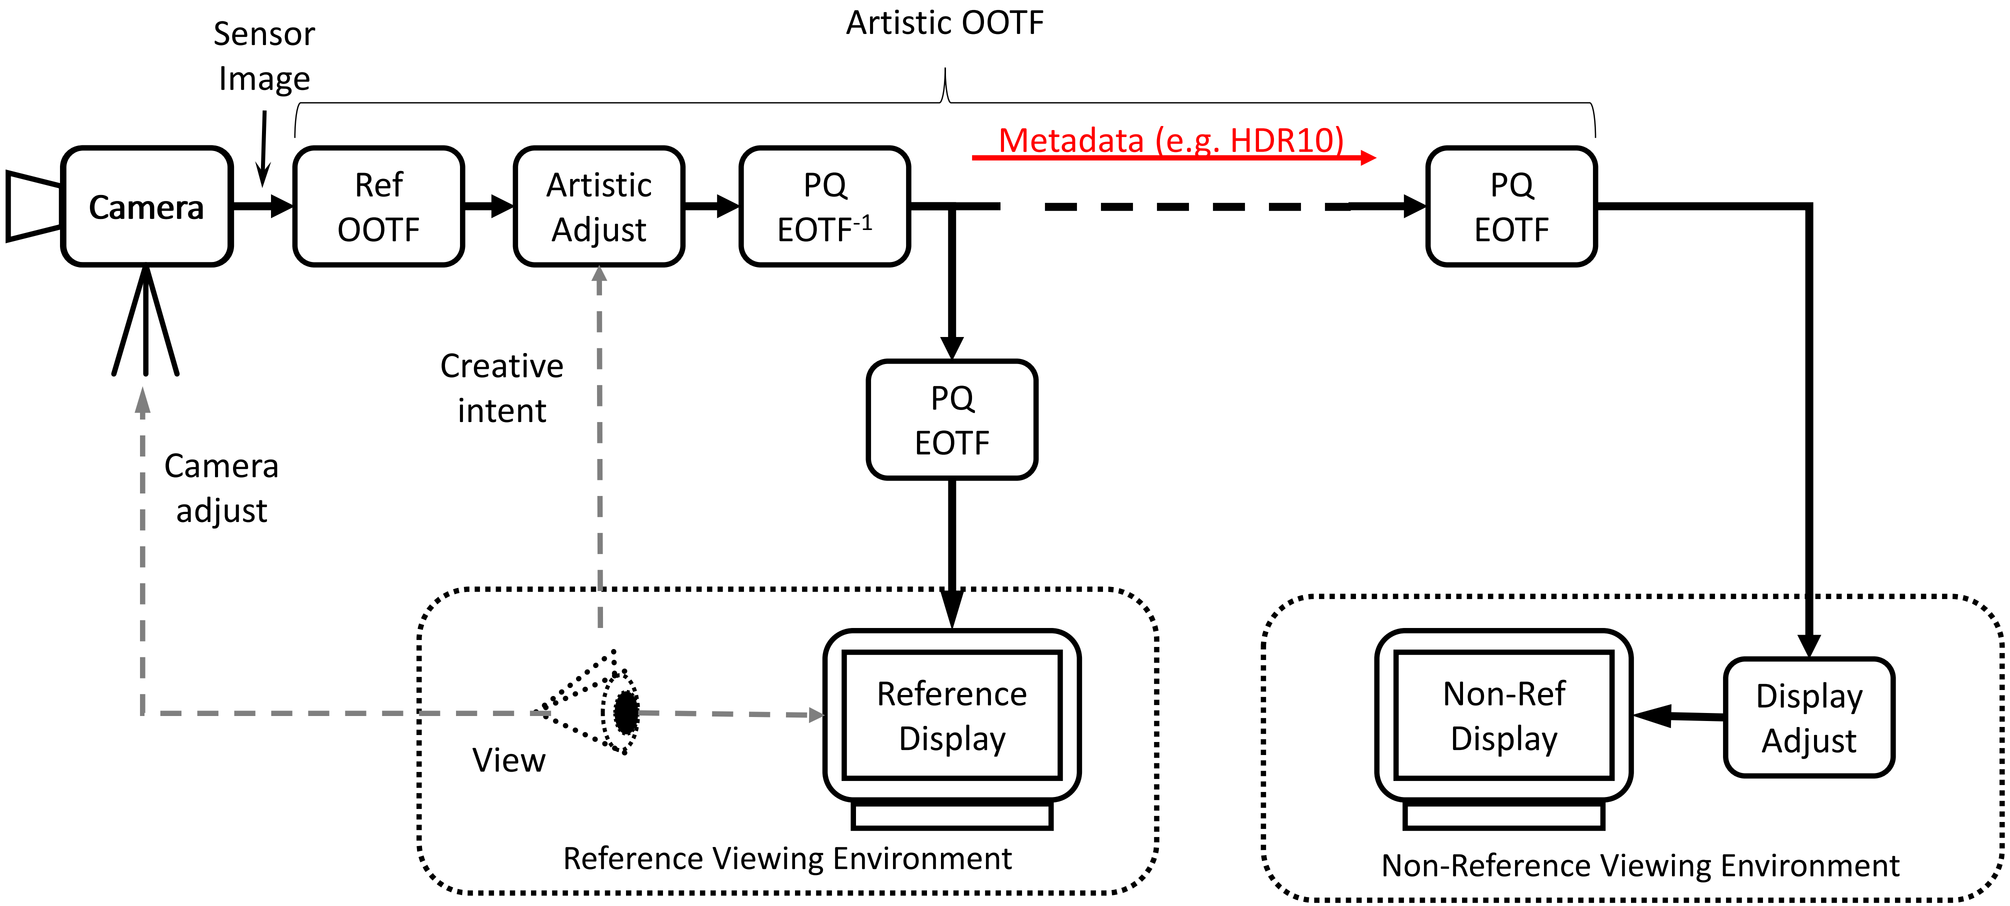

Based on ITU-R BT.2390-4 FIGURE 12

PQ reference broadcast model based on Report ITU-R BT.2390-4 (04/2018) and Recommendation ITU-R BT.2100-2 (07/2018)

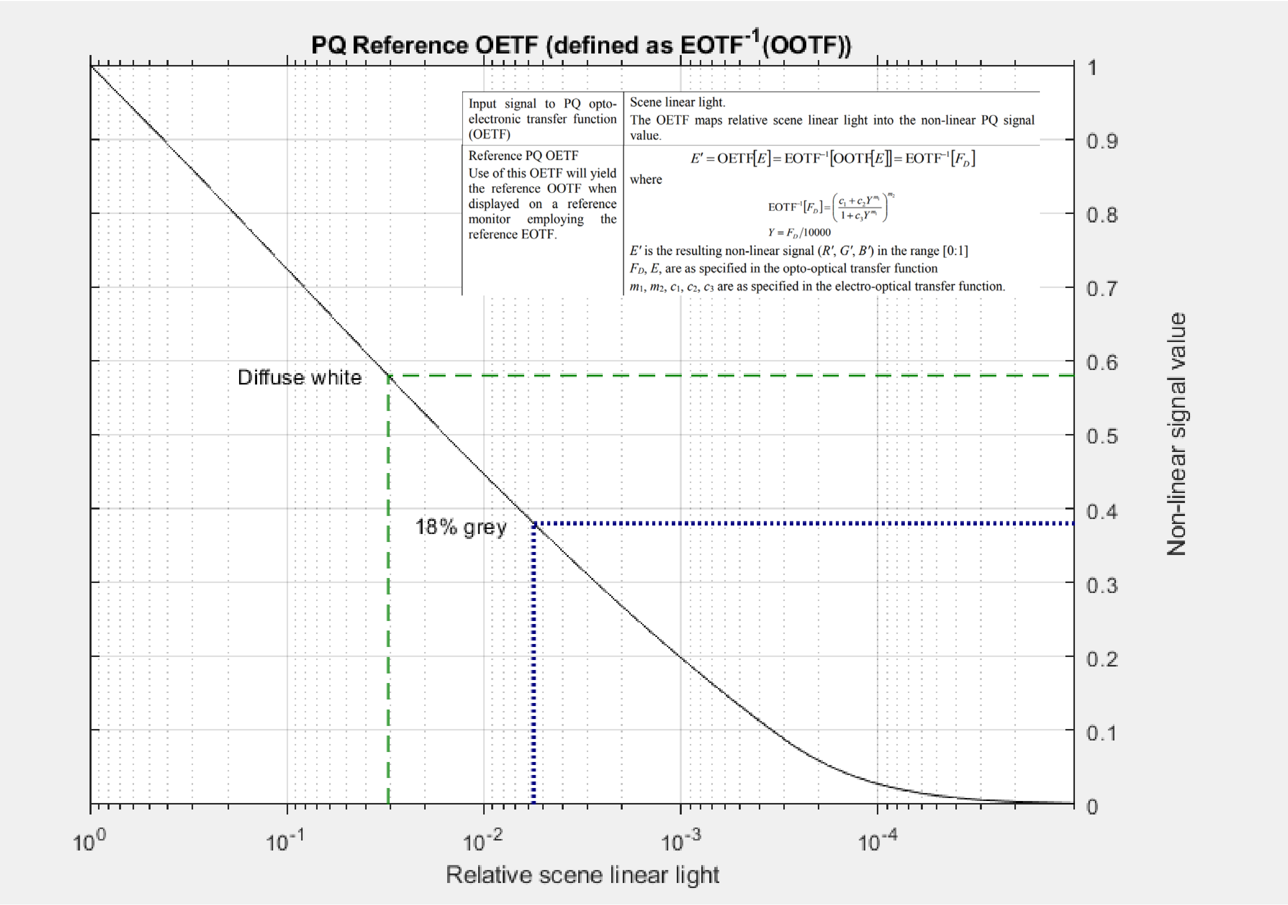

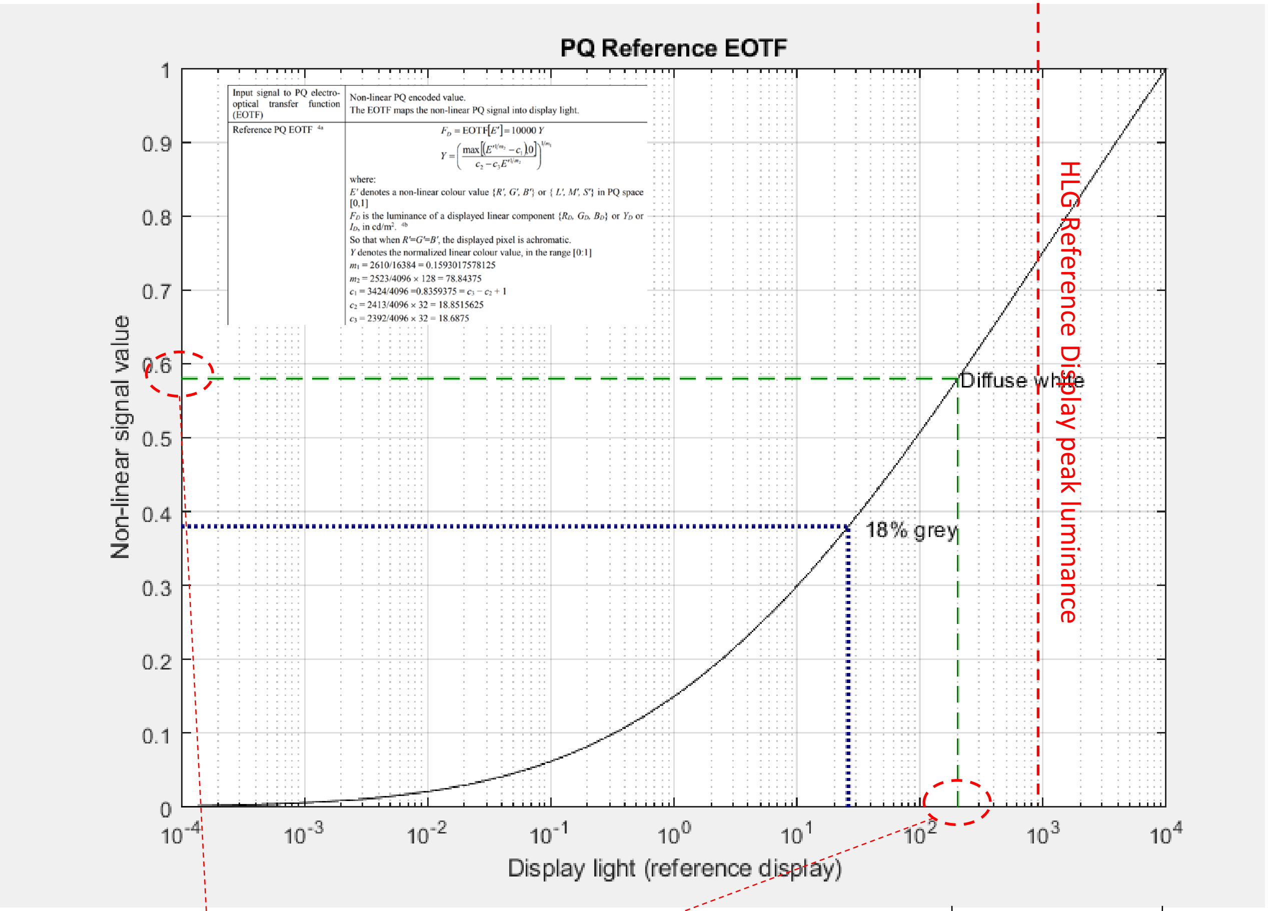

Peak white is 32.73 x diffuse white

Peak white is 49.59 x diffuse white

Nominal signal levels for PQ and HLG production

Reflectance Object or Reference (Luminance Factor, %)[3]	Nominal Luminance, cd/m² (PQ & 1000 cd/m² HLG)	Nominal Signal Level %PQ	Nominal Signal Level %HLG
Grey Card (18%)	26	38	38
Greyscale Chart Max (83%)	162	56	71
Greyscale Chart Max (90%)	179	57	73
Reference Level: HDR Reference White (100%) also diffuse white and Graphics White	203	58	75

TABLE 1 from Report ITU-R BT.2408-1

NOTE LOG SCALE

Problem with PQ and ICC Colour Management

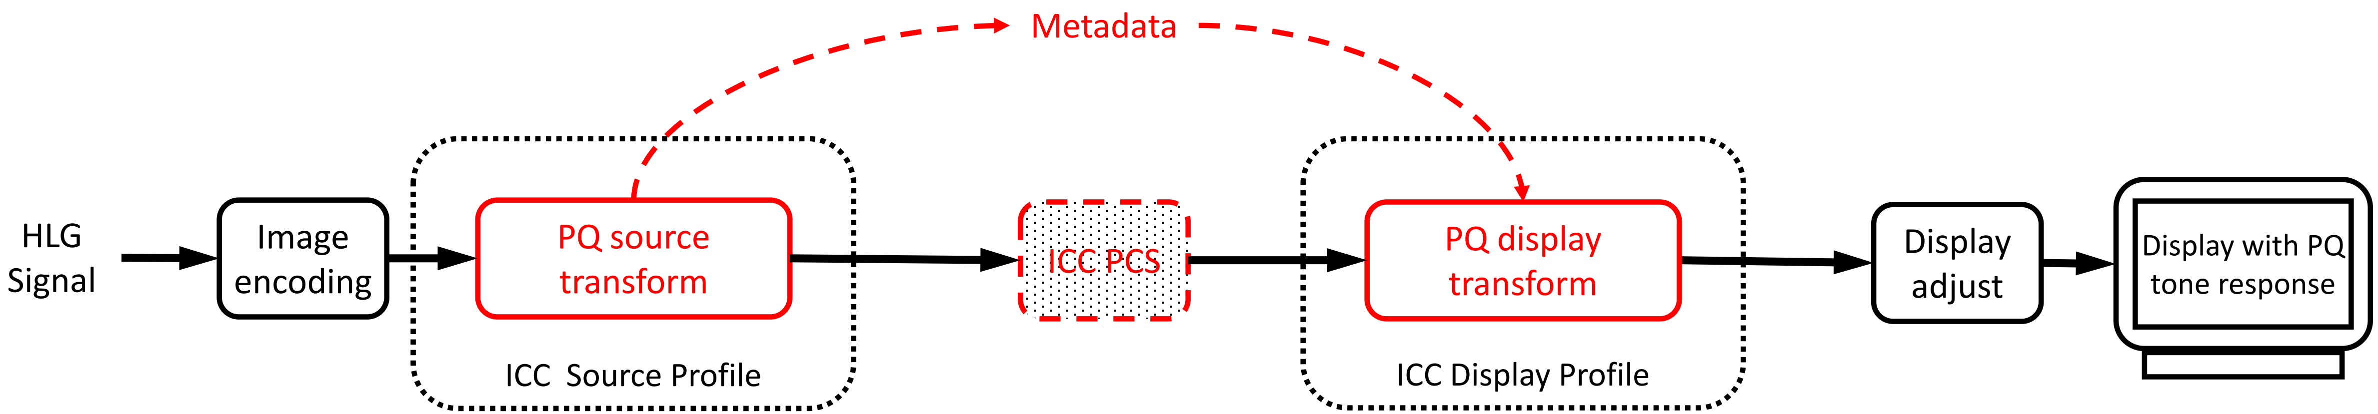

The PQ display profile transform depends on metadata from the source transform.
Metadata includes:

- Maximum Frame–Average Light Level (MaxFALL)
- Maximum Content Light Level (MaxCLL)
- Colour gamut

Extension to HLG model to incorporate PQ (HDR10)

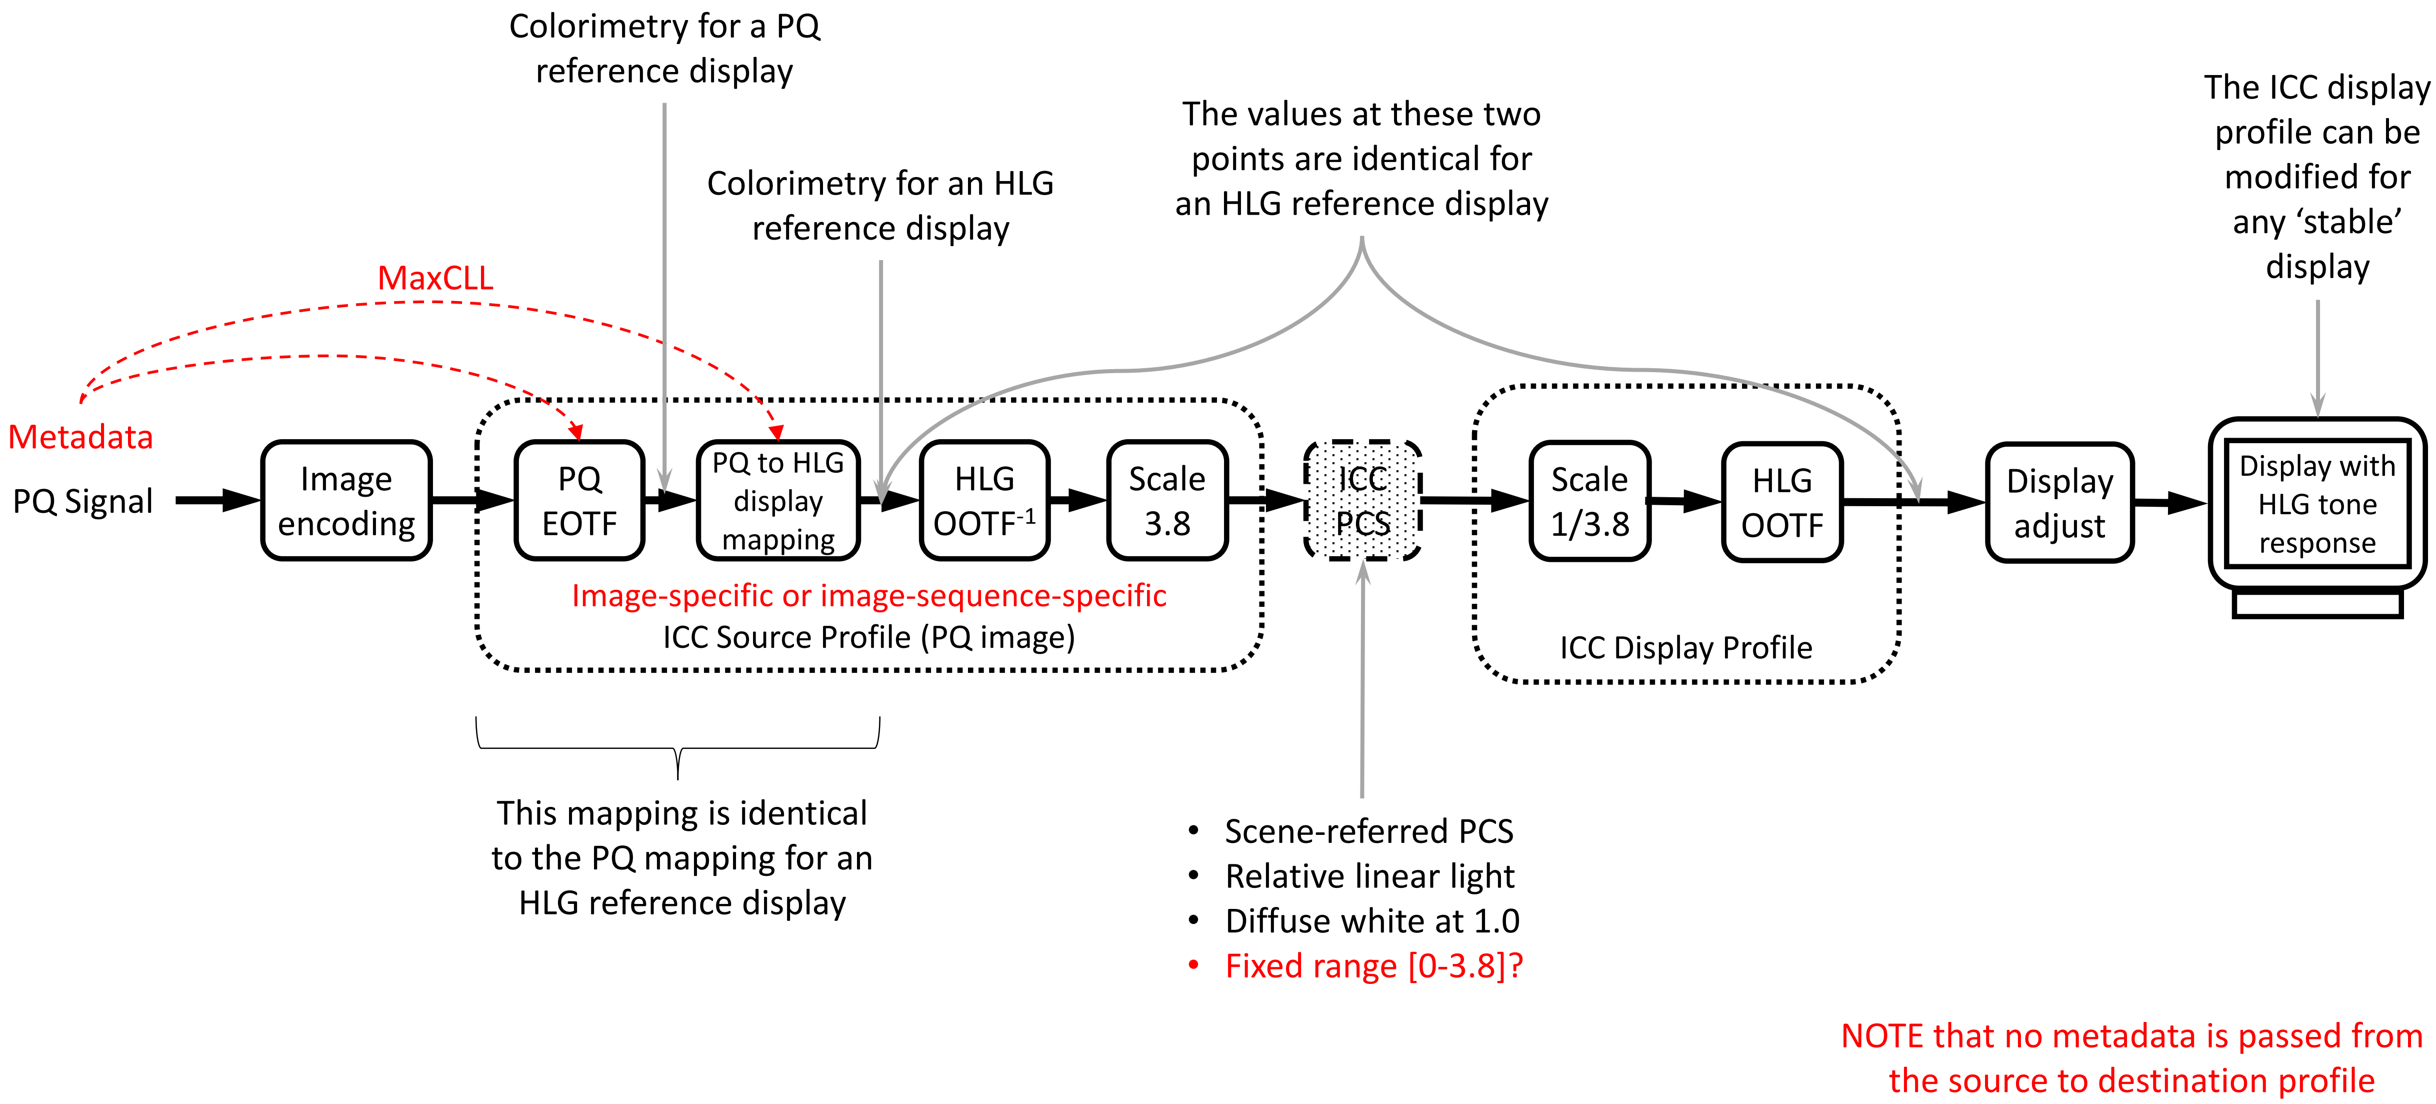

Possible ICC HDR support for PQ based on Report ITU-R BT.2390-4 (04/2018) and Recommendation ITU-R BT.2100-2 (07/2018)

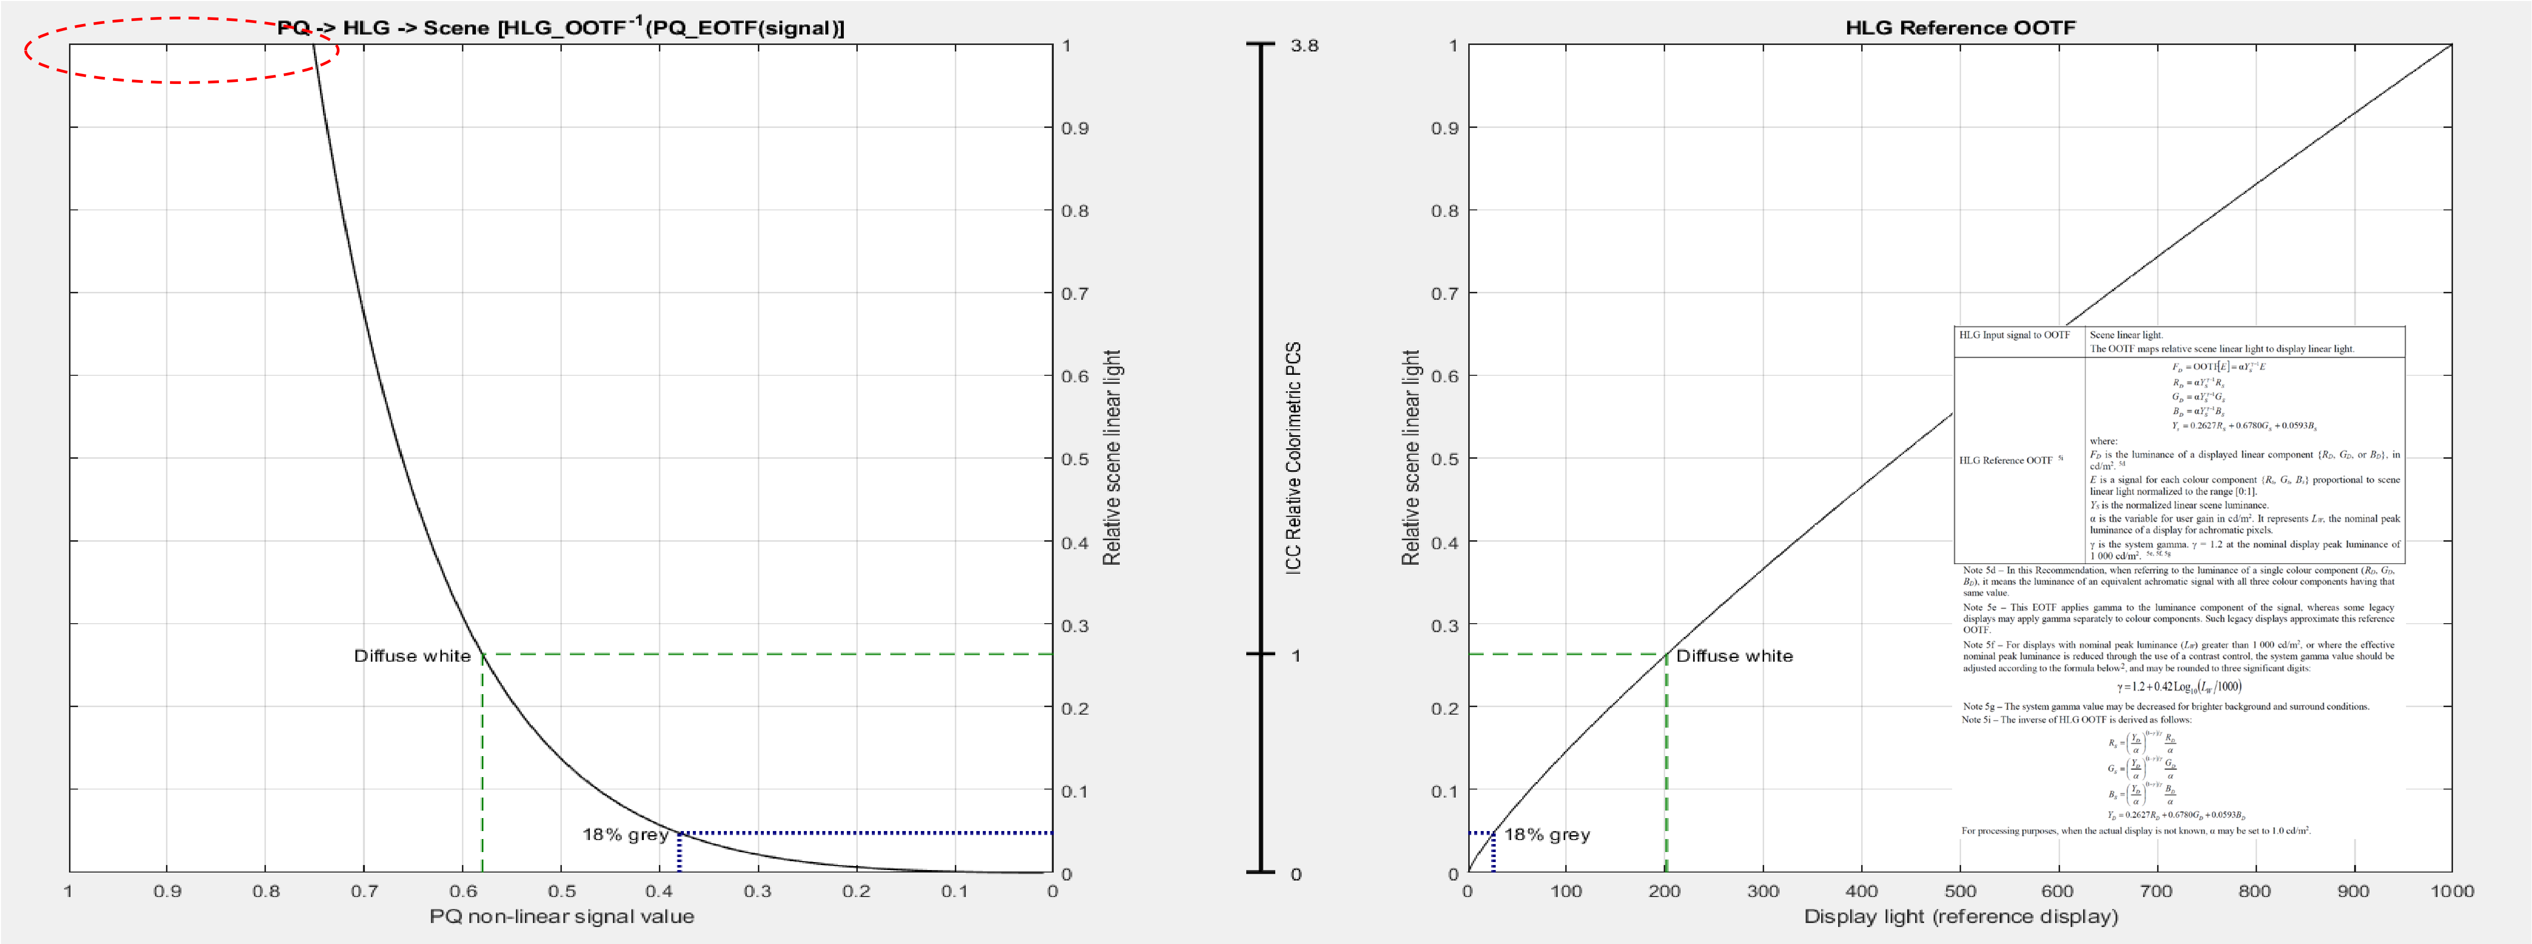

Note that PQ highlights corresponding to signal values greater than 0.75 will be clipped by this transform

It may be desirable to map these in some way to avoid clipping

The display profile is the same for HLG and PQ source data

PQ broadcast model for HLG reference display based on Report ITU-R BT.2390-5, Recommendation ITU-R BT.2100-2

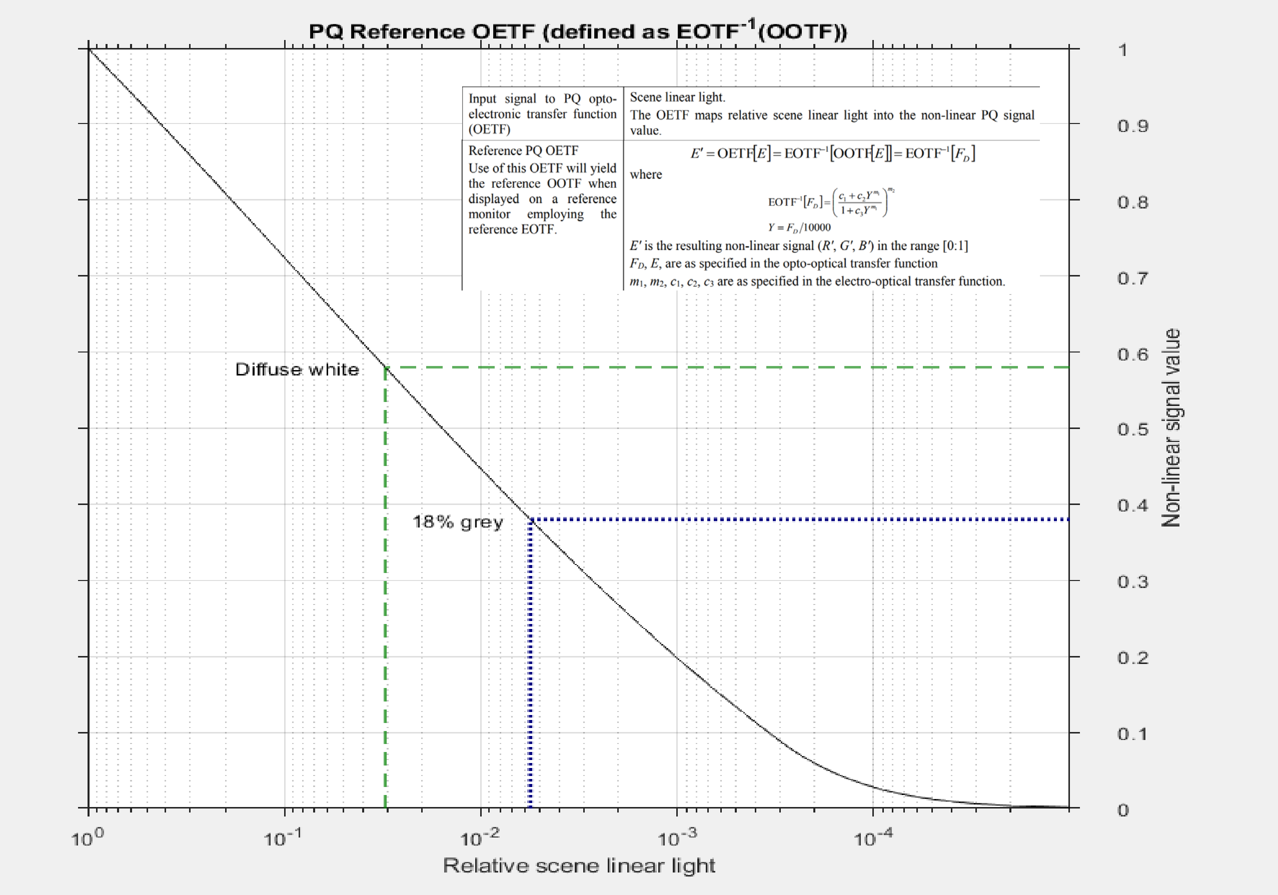

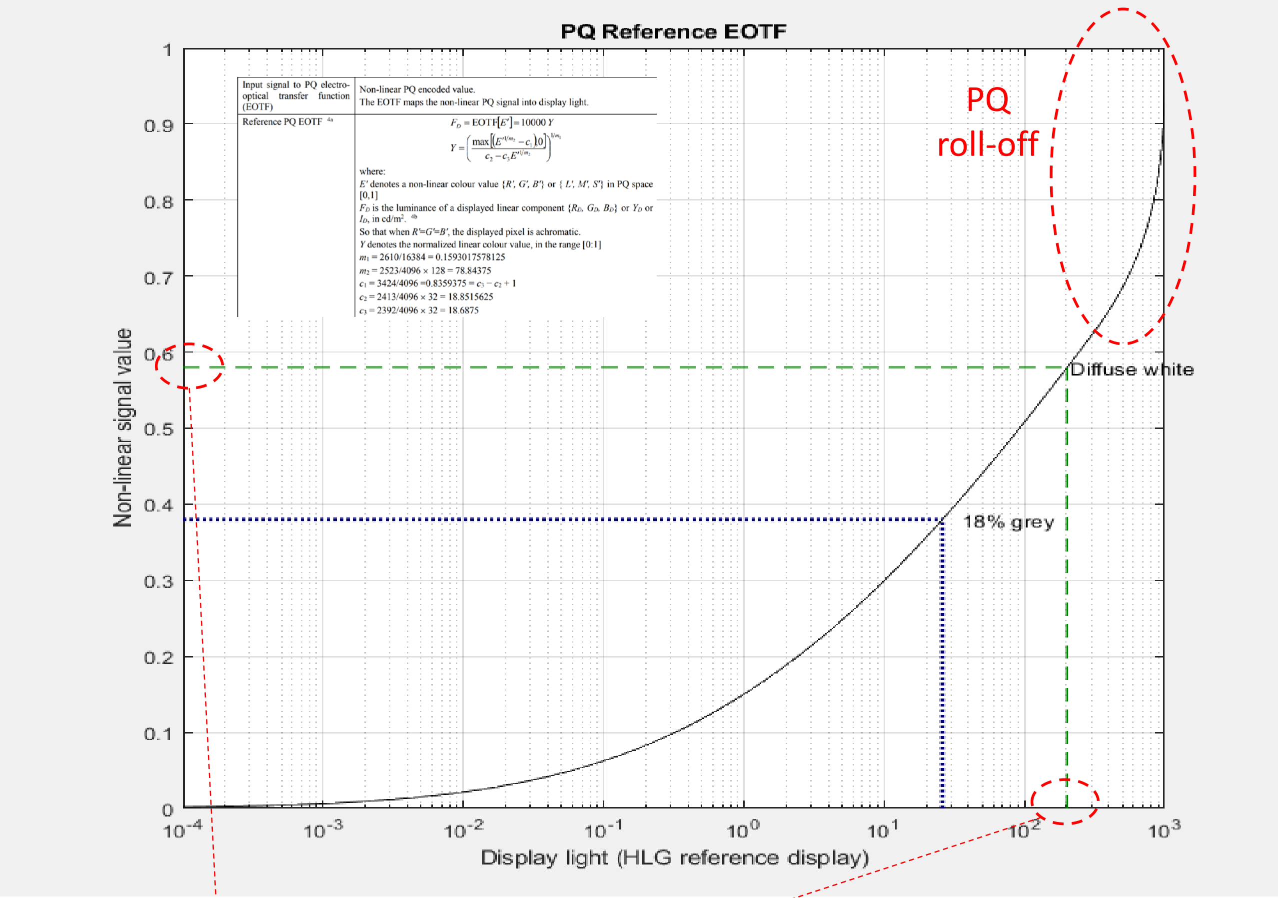

Nominal signal levels for PQ and HLG production

Reflectance Object or Reference (Luminance Factor, %)[3]	Nominal Luminance, cd/m² (PQ & 1000 cd/m² HLG)	Nominal Signal Level	
		%PQ	%HLG
Grey Card (18%)	26	38	38
Greyscale Chart Max (83%)	162	56	71
Greyscale Chart Max (90%)	179	57	73
Reference Level: HDR Reference White (100%) also diffuse white and Graphics White	203	58	75

TABLE 1 from Report ITU-R BT.2408-1

Possible ICC HDR support for PQ based on based on Report ITU-R BT.2390-5, Recommendation ITU-R BT.2100-2

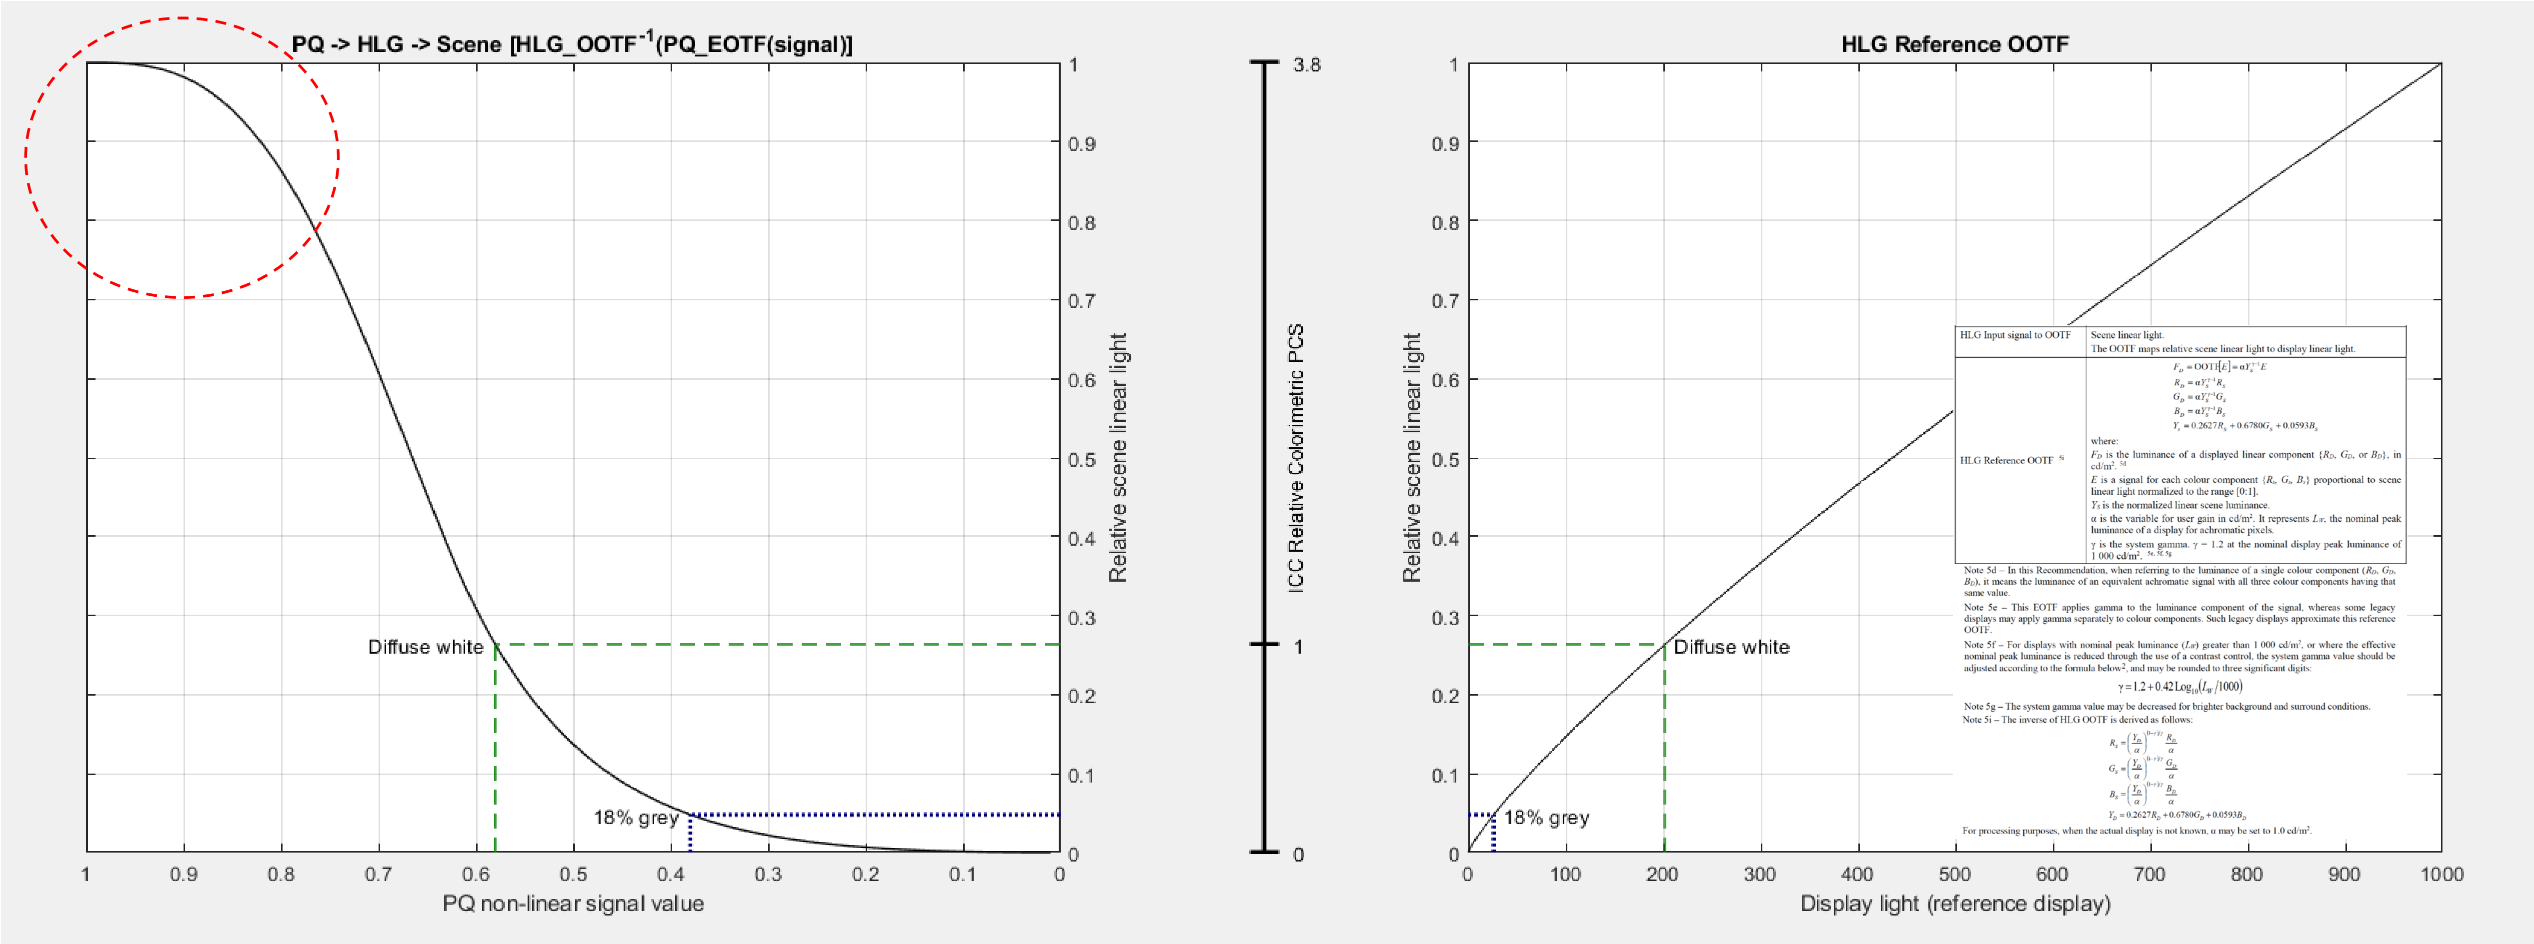

PQ highlights corresponding to signal values greater than 0.75 remapped using PQ EETF to avoid clipping of highlights

The display profile is the same for HLG and PQ source data

Some observations

- HDR displays are typically adjusted dependent on ambient conditions and there is generally no way to know that they have been adjusted
- Unless the dynamic range of the HDR display is the same as the HDR reference display some form of tone reshaping is desirable in order to avoid clipping
- This requires a (minor) rethinking of what is meant by ICC Relative Colorimetric intent
- The ICC Specification currently uses a different model for encoding of overrange values which (in my opinion) makes it harder to produce reliably good results
- In many CMMs, PCS values greater that 1 are clipped; add a rule (as is implemented in the Adobe Color Engine) that if both source and destination profiles use the extended PCS range no clipping is performed

Implementation notes (from Adobe)

Limited precision in mAB mBA matrices limits max saturation

- The matrix encoding is s15.16. The 16 bit fraction results in quantization errors up to ½*1/65536 in roundtrip conversions that combine the mAB and mBA matrices. This quantization shows up as channel cross talk at very high saturation levels. At >99.5 % saturation in encoding space, this quantization noise dominates the signal. Some workarounds:
    1. Use the matrix TRC encoding, and let the CMM create a higher-precision inverse of the AB matrix.
    2. Tweak the mAB and mBA matrix pair, reducing the quantization error by a factor of 10.
    3. If the CMM supports intermediate overrange values, scale up the matrix values and scale down the results using M curves

Curv tag precision not enough for HDR

- The curv tag on its own supports a dynamic range of at most 16 stops (measured in the PCS XYZ space). The dynamic range can be increased by pairing the curv tag with a gamma curve. The curv tag goes into the A curve, and the gamma curve goes into the M curve. For example, a gamma of 2 (or ½) can increase the dynamic range to 32 stops.

16-bit Integer PCS encoding not enough for HDR or 16-bit images

- Some CMMs seem to use 16-bit unsigned integers as the internal PCS encoding. 16-bit unsigned integers support at most 16 stops, such as Adobe RGB values down to 7/1023. For a gamma-encoded color space, the PCS needs to support (bitdepth*gamma) stops or bits. Current HDR needs up to 24 stops. Floats or 32-bit integers support at least 32 stops.

Slope limits hurt shadow details

- HDR shadow values go down to 0.000005 or less. Some CMMs use slope limits for values < 1/100. These slope limits can cause conversion errors.

What next?

- High Dynamic Range imaging is a major step in colour imaging technology, it seems that the ICC can provide a useful forum for discussion of this topic and perhaps contribute towards a solution
- Can we engage stakeholders (OS vendors, Display manufacturers, VESA, W3C, Content providers, Dolby,…)

end of presentation

Reference viewing environment for critical viewing of HDR programme material

Parameter	Values
Surround and periphery [3a]	Neutral grey at D65
Luminance of surround	5 cd/m^2
Luminance of periphery	$\leq$ 5 cd/m^2
Ambient lighting	Avoid light falling on the screen
Viewing distance [3b]	For 1 920 × 1 080 format: 3.2 picture heights For 3 840 × 2 160 format: 1.6 to 3.2 picture heights For 7 680 × 4 320 format: 0.8 to 3.2 picture heights
Peak luminance of display [3c]	$\geq$ 1 000 cd/m^2
Minimum luminance of display (black level) [3d]	$\leq$ 0.005 cd/m^2

Note 3a – “Surround” is the area surrounding a display that can affect the adaptation of the eye, typically the wall or curtain behind the display; “periphery” is the remaining environment outside of the surround.

Note 3b – When picture evaluation involves resolution, the lower value of viewing distance should be used. When resolution is not being evaluated, any viewing distance in the indicated range may be used.

Note 3c – This is not to imply this level of luminance must be achieved for full screen white, rather for small area highlights.

Note 3d – For PQ in a non-reference viewing environment, or for HLG (in any viewing environment), the black level should be adjusted using the PLUGE test signal and procedure specified in Recommendation ITU-R BT.814.

TABLE 3 from Recommendation ITU-R BT.2100-2 (07/2018)

Explanation from Tim Borer (BBC R&D)

- It IS reasonable to describe HLG as relative scene luminance
  - There has been a lot of debate by proponents of PQ (which is absolute/non-relative display referred encoding) that this is not really so. Nevertheless we designed HLG precisely as relative scene luminance (just like Rec.709).
- I have often heard it said that there is no standard Rec.709 production (because camera operators, it is claimed, universally adjust their cameras)
  - The claim is that by tweaking the camera the picture somehow becomes display referred. Even if it were true that cameras are always adjusted (not so) this would not make the signal display referred. If you doubt this simply look at the dimensions of the signal. HLG is dimensionless (a relative signal) and PQ has dimensions of candelas per square metre (nits)
  - All that adjusting a 709/HLG camera does is to produce an “artistic” modification to the signal. The signal still represents relative scene referred light, just not the actual scene (but, rather, one that the producer wished had existed). Adjusting the camera does not convert a dimensionless signal into a dimensional one.
  - In fact a great deal of live television is produced using a standard Rec709 OETF. This includes almost all live sport (especially soccer), live light entertainment, and news. This encompasses a large part of broadcast television output. In sport it is often a contractual obligation to use the Rec 709 OETF. In other instances the producers often do not like knees because, as typically implemented, then can distort flesh tones. A further consideration it that in a multicamera shoot it is essential to match all the cameras. This is difficult if they don’t use a standard Rec 709 setting (often cameras may have different firmware versions, which means that setting up a camera does not necessarily mean the same thing even on the same model of camera)
- It is not the case that “the camera’s linearity response is tweaked by the operator for various reasons”
  - This is not really viable in a multicamera live shoot, and the shaders don’t have time to do it live
  - Similarly for live production the gamut is standard Rec 709 (if necessary clipped from the wider camera taking gamut). This is necessary to ensure consistent colours for sporting strips both between cameras and at different venues and different games. It is not unusual for sporting strips to be outside 709, and it is important that such colours are treated in a consistent way, so gamut mapping, other than simple clipping, is not viable (bear in mind that footage from different games may often be shown as part of the commentary). So, to emphasise, the fact is that a great deal of television IS produced using the standard Rec 709 OETF using standard 709 gamut (without gamut mapping)

Explanation from Tim Borer (BBC R&D)

- Scene referred conversions need to be used appropriately
  - When used appropriately they do not create significant colour shifts. Scene referred conversion should be used when matching camera outputs (as opposed to matching the picture that is seen on the display – they are not the same).
  - Display referred conversions are used to ensure that the displayed image is the same. There are different, distinct, use cases for these two types of conversion. It is a mistake to assume that all conversions should be display referred. As an example consider the recent coverage of the Royal Wedding. This was shot using a mixture of HLG HDR cameras and HD cameras (using Rec 709). The production architecture was that shown in ITU-R Report BT.2408-1 2018 (note the -1 version) (freely available at https://www.itu.int/pub/R-REP-BT.2408-1-2018), figure 4, page 14. In this workflow you will note that there are many scene referred Rec 709 to HLG conversions. This is so that pictures can be shaded using standard Rec 709 monitors (which is a requirement when the majority of viewers are watching in 709). Note, in particular, that the final Rec 709 SDR output is converted from the HLG signal using a scene referred conversion. We estimate that the Royal Wedding was viewed by 1.8 billion viewers. The international feed was derived from the HLG signal as described, using a scene referred conversion. The colour was NOT distorted using this conversion (and, clearly, the producers would not have allowed distortion for such a prestigious broadcast). On the other hand if one were producing primarily for an HDR audience then you would use the alternative architecture in BT.2408 (fig 3, page 12). Here shading is performed primarily on the HDR monitor (with an SDR monitor fed by a DISPLAY referred conversion, so that the shader can check that nothing untoward is happening on the SDR signal). Note that in a joint HDR SDR production you can give primacy to either shading in HDR (for a majority HDR audience) or to shading in SDR (for a mainly SDR audience – the current situation). But you cannot prioritise both. We have found that shading in SDR gives very good quality HDR as well as SDR. NHK have in mind producing for a primarily HDR audience and, therefore, they favour the production workflow of figure 3. To summarise scene referred conversion do not produce significant colour shifts when they are used properly. Indeed without using scene referred conversions it is not possible to match the look of HDR and SDR output when shading in SDR.
- There are few HDR broadcasts in Europe yet though a number are in the pipeline
  - To the best of my knowledge all these broadcasts will be using HLG
  - Similarly broadcasters that produce live content (such as sport) in the US also favour HLG
  - The BBC has made our Planet Earth II series available in UHD HDR HLG on catch up (OTT) television, and are showing some of the Soccer World Cup matches is HLG HDR OTT. Most OTT movie distribution currently uses PQ (HDR10). This is possible because of the non-live workflow
  - I am unaware of live sporting events broadcast using PQ
  - YouTube distribute in both HLG and HDR10. Since production is easier in HLG (particularly with pro-summer equipment) most of the user generated HDR content on YouTube is HLG.

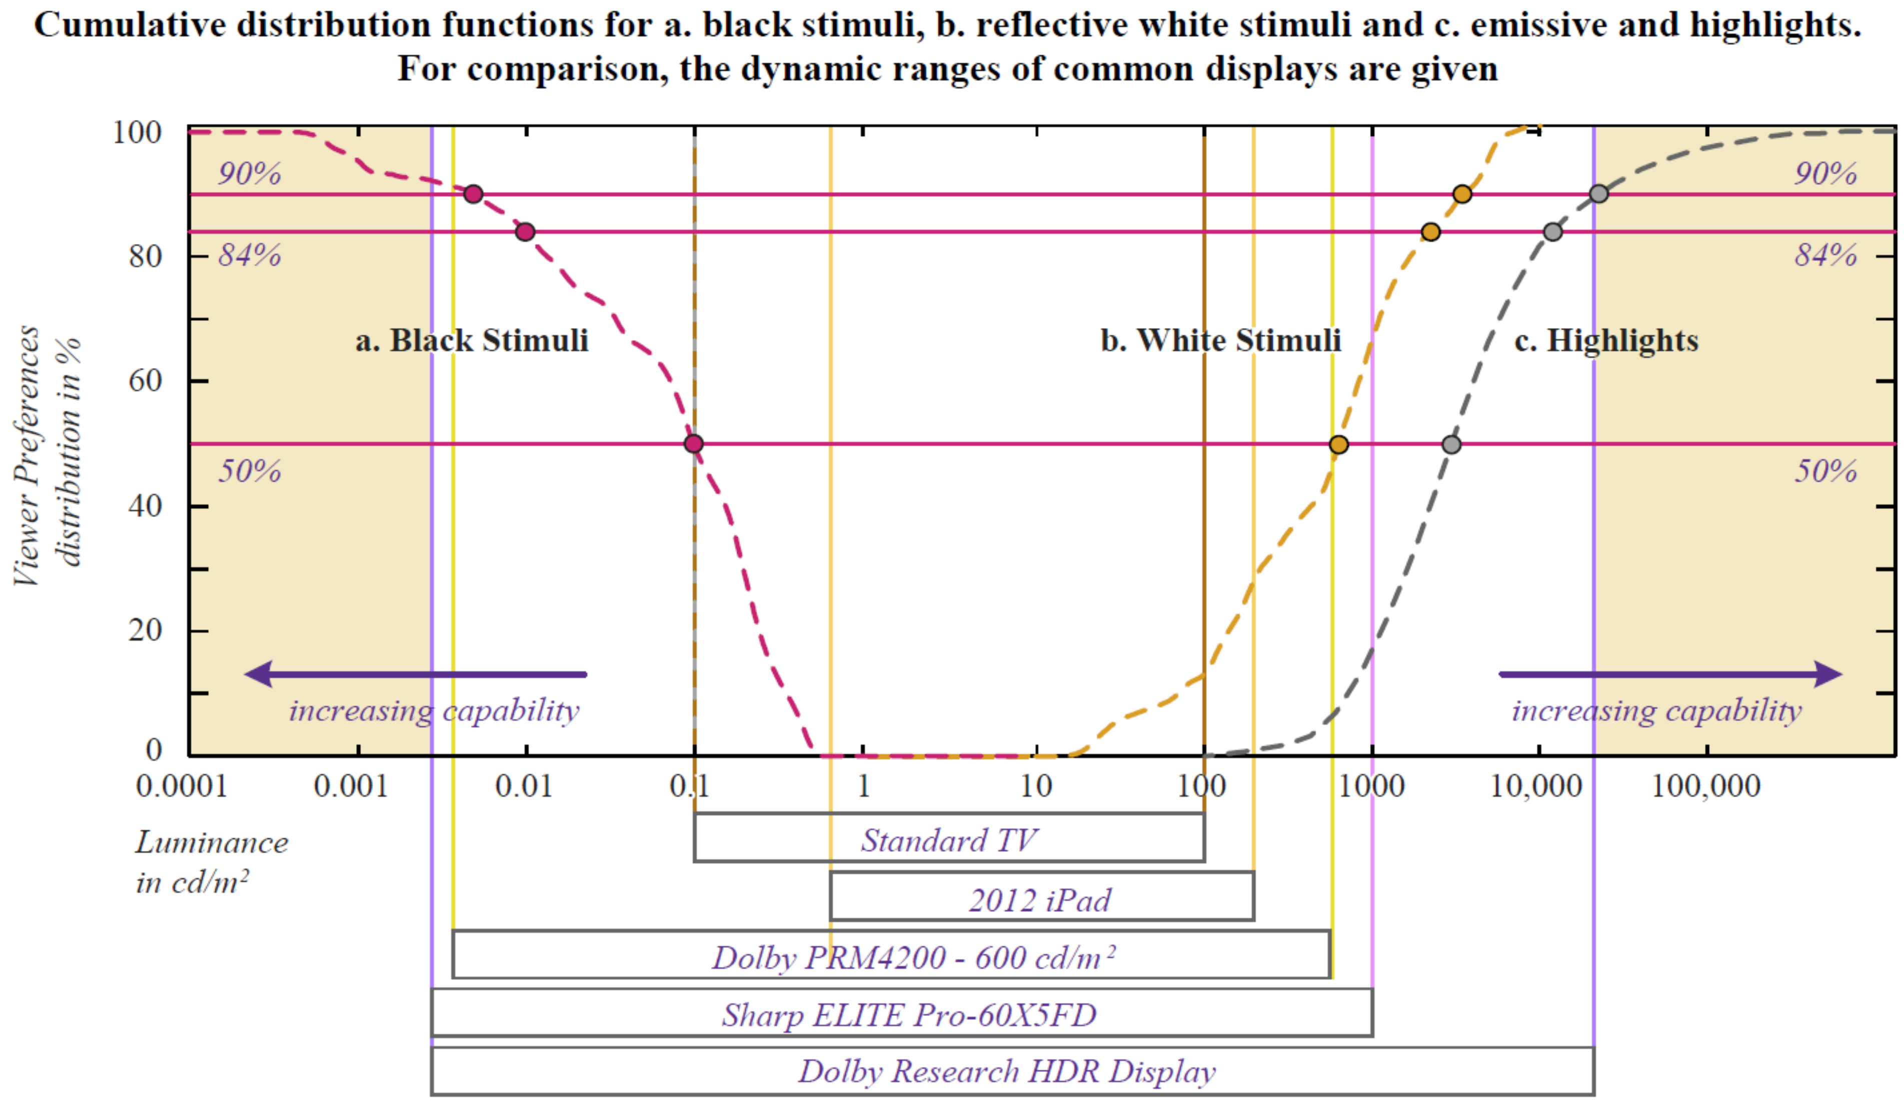

ITU-R BT.2390-4 FIGURE 3