Color Threshold Functions: Application of Contrast Sensitivity Functions in Standard and High Dynamic Range Color Spaces

Minjung Kim, Maryam Azimi, and Rafał K. Mantiuk
Department of Computer Science and Technology, University of Cambridge

Abstract

Contrast sensitivity functions (CSFs) describe the smallest visible contrast across a range of stimulus and viewing parameters. CSFs are useful for imaging and video applications, as contrast thresholds describe the maximum of color reproduction error that is invisible to the human observer. However, existing CSFs are limited. First, they are typically only defined for achromatic contrast. Second, even when they are defined for chromatic contrast, the thresholds are described along the cardinal dimensions of linear opponent color spaces, and therefore are difficult to relate to the dimensions of more commonly used color spaces, such as sRGB or CIE $L^*a^*b^*$. Here, we adapt a recently proposed CSF to what we call color threshold functions (CTFs), which describe thresholds for color differences in more commonly used color spaces. We include color spaces with standard dynamic range gamut (sRGB, YC_bC_r, CIE $L^*a^*b^*$, CIE $L^*u^*v^*$) and high dynamic range gamut (PQ-RGB, PQ-YC_bC_r and IC_TC_P). Using CTFs, we analyze these color spaces in terms of coding efficiency and contrast threshold uniformity.

Introduction and Background

Contrast thresholds describe the minimum difference in luminance or chromaticity that a human observer can detect. Contrast thresholds vary with image and viewing parameters, such as spatial frequency [1], luminance [2], color [3], and the size of the stimulus [4]. The dependence of contrast thresholds to such parameters is described by contrast sensitivity functions (CSFs).

Having an accurate CSF is important for image and video coding applications, as contrast thresholds provide a limit to the amount of color reproduction error that is noticeable to a human observer. For example, CSFs were used to create transfer functions for encoding high dynamic range color values [5, 6] in a perceptually uniform manner. In particular, Perceptual Quantizer (PQ) is a non-linear function based on Barten's CSF [7], used to code High Dynamic Range (HDR) content [6], akin to gamma encoding used in Standard Dynamic Range (SDR) [8]. For fixed bit-depth, PQ assigns more code-words to luminance levels with lower thresholds and fewer code-words to levels with high thresholds, thus maximizing visual quality while efficiently allocating code-words.

However, many existing CSFs only describe achromatic contrast, neglecting the detection thresholds in chromatic directions. This is true even for Barten's CSF, meaning that PQ is only appropriate for coding luminance, even though the industry standard is to use PQ for all color channels [9, 10]. In addition, existing chromatic CSFs [11, 12] are usually reported in Derrington-Krauskopf-Lennie (DKL) color space [13], a linear color opponent space that is physiologically relevant but has limited applica-

tion outside of vision science. However, it is not obvious as to how contrast thresholds in DKL would translate to thresholds in other color spaces across their color components due to non-linearities such as PQ encoding.

In this work, we adapt our spatio-chromatic CSF[1] [14] to predict *color threshold functions* (CTFs). CTFs describe detection thresholds in color spaces that are more commonly used in imaging, video, and color science applications, such as sRGB and YC_bC_r. The spatio-chromatic CSF model from [14] can predict detection thresholds for any point in the color space and for any chromatic and achromatic modulation. In addition, the CSF was fitted to the data that span 0.0002 cd/m^2 (scotopic) to 10,000 cd/m^2 (photopic), which makes it appropriate for predicting thresholds in HDR color spaces. Using this CSF, we numerically solve for detection thresholds in any non-linear color space. Our work lends insight into the coding efficiency and the uniformity of contrast thresholds in different color spaces.

A Device-Independent CSF

The spatio-chromatic CSF from [14] is the basis for the CTFs presented in this work. The CSF was developed to account for contrast threshold measurements from 0.125 cycles per degree (cpd) to 32 cpd, from 0.0002 cd/m^2 to 10,000 cd/m^2, and for different hues [16, 17, 18, 15, 19, 20]. A critical feature of this CSF is that it can accurately describe contrast thresholds for any luminance and chromaticity combination, meaning we can map contrast thresholds from the native color space of the CSF onto color coordinates of other color spaces.

In Fig. 1, we show the CSF in DKL color space, a linear transformation of LMS color space:

$$\Delta DKL = \begin{bmatrix} 1 & 1 & 0 \\ 1 & -\dfrac{L_0}{M_0} & 0 \\ -1 & -1 & \dfrac{L_0 + M_0}{S_0} \end{bmatrix} \begin{bmatrix} \Delta L \\ \Delta M \\ \Delta S \end{bmatrix} \quad (1)$$

(L_0, M_0, S_0) is the white point in LMS color space and $(\Delta L, \Delta M, \Delta S)$ is the color coordinate with respect to the white point. For the CSFs in Fig. 1 and the CTFs derived from them, we used CIE 2006 cone fundamentals [21] and assumed D65 white point. ΔDKL is the color coordinate in DKL, representing position along achromatic, red-green, and yellow-violet axes. The

---

[1]The code and more details on the spatio-chromatic CSF can be found at `https://www.cl.cam.ac.uk/research/rainbow/projects/hdr-csf/`.
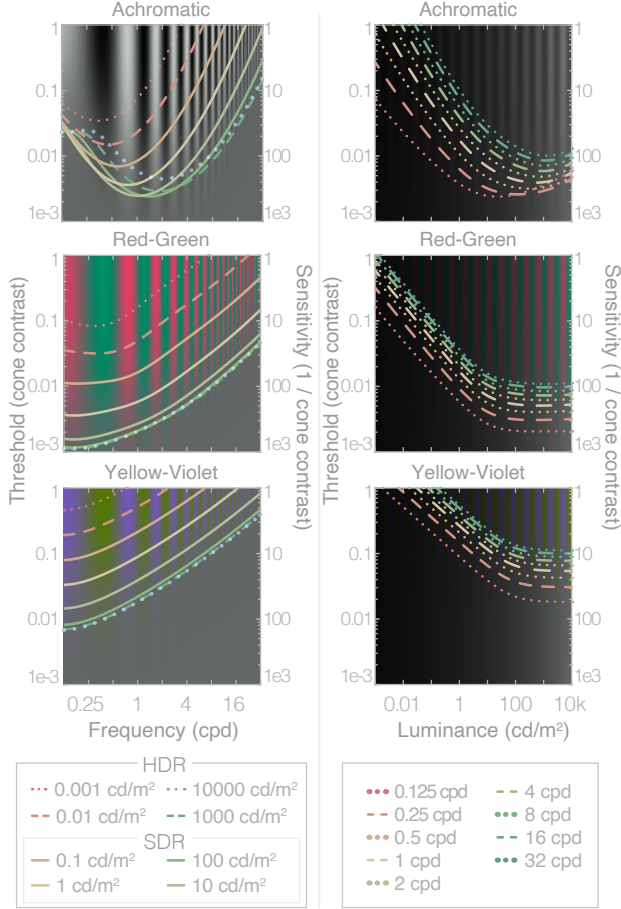
Figure 1: Spatio-chromatic contrast threshold functions (CSFs) along the cardinal directions of the DKL color space [14, 15]. Left column: CSFs as a function of frequency. Right column: CSFs as a function of luminance. Left y-axis: Contrast threshold. Right y-axis: Contrast sensitivity.

contrast thresholds are expressed using cone contrast, a vector length in LMS space:

$$C_{\text{cone}} = \sqrt{\frac{1}{3}\left[\left(\frac{\Delta L}{L_0}\right)^2 + \left(\frac{\Delta M}{M_0}\right)^2 + \left(\frac{\Delta S}{S_0}\right)^2\right]} \qquad (2)$$

The scale factor $1/3$ ensures that $C_{\text{cone}} \in [0,1]$. At threshold, the probability of detection is $P_{\text{det}} = 0.5$. Using DKL space with cone contrast allows a device-independent representation of the CSF.

The detection thresholds in Fig. 1, as well as for all CTFs, are for Gabor functions extending a $40\,° \times 40\,°$ area. This is approximately equal to viewing a 75-inch TV from a distance of 260 cm ($3 \times$ the height of the display) and is the same assumption used to derive PQ [6].

CTFs in Commonly Used Color Spaces

The goal is to find detection thresholds in any color space, including those with non-linear transfer functions. To achieve that, we start from the D65 white point of the target color space, then search along each axis of the space until we find a color incre-

ment Δc that yields the detection threshold. For example, consider sRGB. We start at $(R_{\text{D65}}, G_{\text{D65}}, B_{\text{D65}})$ and search for ΔR such that $(R_{\text{D65}} + \Delta R, G_{\text{D65}}, B_{\text{D65}})$ results in the $P_{\text{det}} = 0.5$, assuming a guess rate of 0. To obtain P_{det}, we convert from sRGB to LMS and query the CSF.

Since Δc is an increment in the non-linear target color space, not a contrast, we refer to our thresholds *increment thresholds* or *color increments thresholds*, and the functions as Contrast Threshold Functions (CTFs). We show increment thresholds ($\Delta c > 0$), but not decrement thresholds ($\Delta c < 0$). This is because the underlying CSF has comparable positive and negative contrast thresholds, resulting in similar increment and decrement thresholds in the target color space.

Below, we report increment thresholds for RGB (SDR and HDR), YC_bC_r (SDR and HDR), CIE L*a*b* and CIE L*u*v* (SDR only), and IC_TC_P (HDR only). For SDR, we used the sRGB non-linearity on ITU-R BT.709 color primaries [22], assuming a display luminance range of $0.1\,\text{cd/m}^2$ to $100\,\text{cd/m}^2$. For HDR, we used PQ as the transfer function on ITU-R BT.2020 color primaries [10], assuming a display luminance range between $0.005\,\text{cd/m}^2$ and $10{,}000\,\text{cd/m}^2$. The values of the increment thresholds between SDR and HDR cannot be directly compared because the SDR color gamut is smaller than that of HDR, meaning an increment threshold of 0.1 corresponds to different amount of physical change for SDR and HDR.

In addition, all plots show the maximum quantization error assuming different bit-depths for each color space (horizontal dotted lines). The maximum quantization error is calculated as:

$$\max \varepsilon_q = 0.5\left(\frac{v_{\max} - v_{\min}}{2^b - 1}\right), \qquad (3)$$

where b is the bit-depth, and $v_{\max}$ and $v_{\min}$ are the maximum and minimum values to be encoded, respectively. Note that $v_{\min}$ can be negative for some color channels (e.g., a^* of CIE L*a*b*). When the error is above the threshold, quantization artifacts are likely visible; when the error is below, the artifacts are likely invisible, however code-words are wasted.

RGB

We show CTFs in RGB with sRGB non-linearity (Fig. 2), PQ encoding (PQ-RGB; Fig. 3), and without any non-linearities at HDR luminance levels (linear RGB; Fig. 4). Without any non-linearities, thresholds rise with increasing luminance. In comparison, the CTFs for sRGB and PQ-RGB are approximately flat, horizontal lines with respect to luminance, meaning that detection thresholds for sRGB and PQ-RGB are more uniform than those of linear RGB as a function of luminance.

The increment thresholds for the green channel are quite low for both SDR and HDR (Fig. 2 and Fig. 3); indeed, industry-standard bit-depths (SDR: 8 bits; HDR: 10 bits) are insufficient for keeping quantization errors below threshold for the green channel. In comparison, the thresholds for the blue channel are high, meaning that an uneven distribution of bits across the channels may be a better coding scheme. For HDR, for example, 12 bits for green, 10 bits for red and 9 bits for blue ensures that all the quantization steps are near or below the detection threshold.
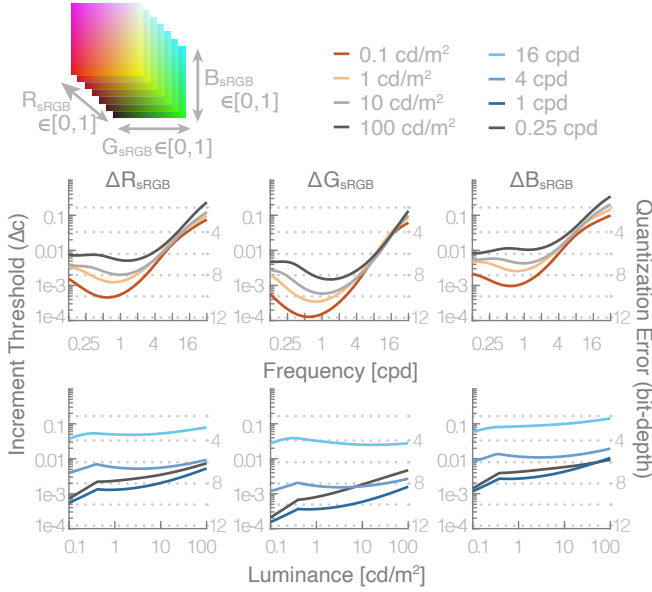
Figure 2: sRGB using ITU-R BT.709 color primaries [22]. Top row: CTF with respect to frequency. Bottom row: CTF with respect to luminance. Left *y*-axis: increment thresholds. Right *y*-axis: quantization error. Dotted grey lines indicate quantization error for different bit-depths. All CTF figures (Figs. 2–9) follow the same formatting convention.
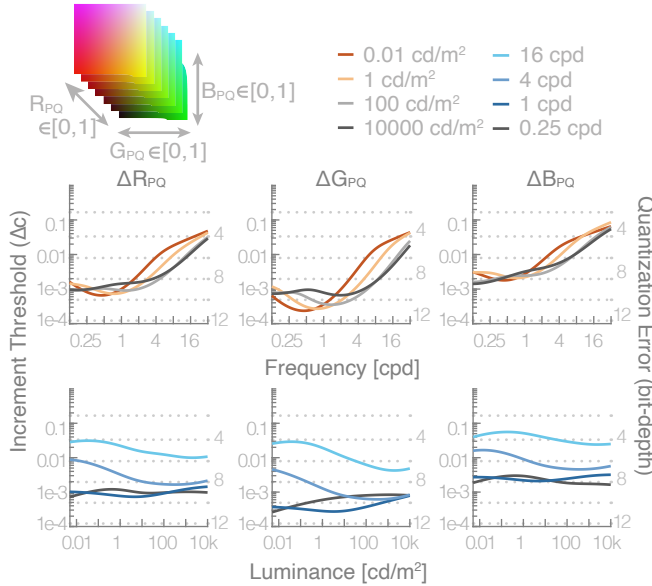

Figure 3: PQ-RGB using ITU-R BT.2020 color primaries [10].

YC_bC_r

YC_bC_r [23] is used for compression standards (e.g., High Efficiency Video Coding [24]), as well as for coding DVD and Blu-ray video. The first channel is called luma (Y) to differentiate it from linear luminance; the other two are blue difference (C_b) and red difference (C_r) from luma. Y is in [0,1]; C_b and C_r are in [-0.5,0.5]. We only consider $\Delta C_b, \Delta C_r > 0$, since $C_b, C_r = 0$ for D65 and we only report color increment thresholds.

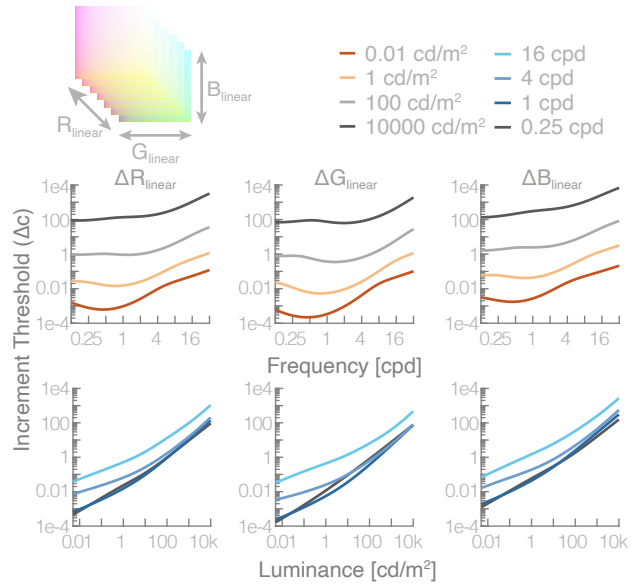For SDR, YC_bC_r is a linear transformation of sRGB; the

Figure 4: Linear RGB using ITU-R BT.2020 color primaries [10]. We do not provide quantization error, since linear RGB is not used for coding.

sRGB non-linearity is inherited by YC_bC_r. The sRGB non-linearity works relatively well for all channels in YC_bC_r, making increment thresholds relatively uniform with respect to luminance (Fig. 5).
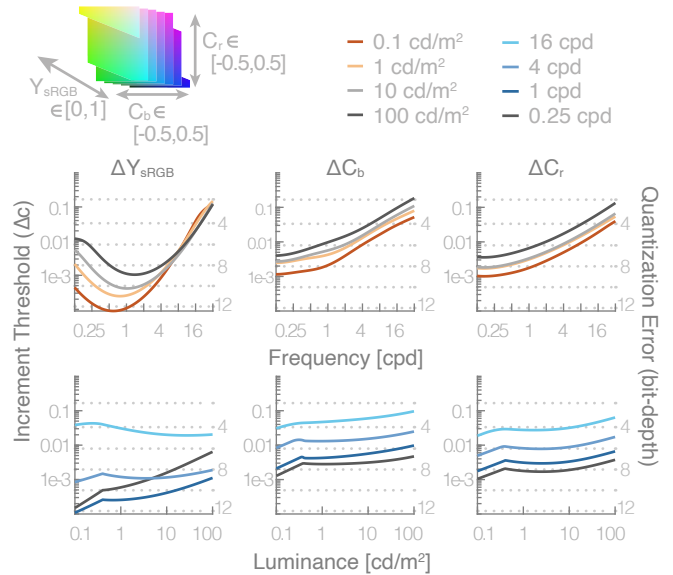

Figure 5: YC_bC_r using ITU-R BT.709 color primaries [22].

For HDR, there are two possible approaches of encoding with PQ [10]: (1) apply PQ on linear luminance directly and acquire luma, and (2) apply PQ on linear RGB, and then acquire luma of YC_bC_r from non-linear RGB. We report the results of the latter method, but have found that the two methods yield comparative CTFs.

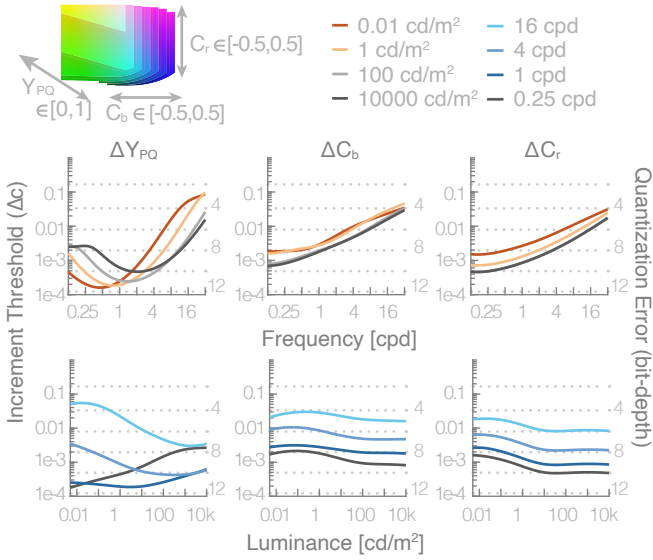Fig. 6 shows the results. Whereas the thresholds in C_b and C_r

Figure 6: PQ-YC$_b$C$_r$ using ITU-R BT.2020 color primaries [10].

Figure 7: CIE L*u*v* using ITU-R BT.709 color primaries [22].

are approximately constant with respect to luminance, thresholds in Y are only constant at 1 cpd. For all three channels, the CTFs are less uniform below 10 cd/m^2 than above, possibly because the CSF underlying PQ was designed for 1 cd/m^2 to 1000 cd/m^2 [6, 7]. The CSF from [14] uses thresholds measured in the entire luminance range of the HDR gamut [15].

The current standards for image and video coding (SDR: 8 bits; HDR: 10 bits) are insufficient to code Y without noticeable error, but are sufficient for the chromatic channels, suggesting the possibility of a coding scheme that better balances code-words across the channels.

CIE L*u*v* and CIE L*a*b*

CIE L*u*v* and CIE L*a*b* are two color spaces adopted by Commission internationale de l'éclairage (CIE) for color matching. These spaces are intended to be perceptually uniform, meaning that one-unit change in color value corresponds to the same amount of perceived color change, for all colors in the space. Both color spaces rely on an assumed white point, are limited to SDR colors, and device-independent. The channels correspond to the lightness (L*), red-green (u* or a*) and blue-yellow (v* or b*) chromatic channels. L* is in [0,100], u* and v* are in [-100,100], a* is in [-500,500] and b* is in [-200,200].

Surprisingly, CTFs in CIE L*u*v* (Fig. 7) and CIE L*a*b* (Fig. 8) are not much different from sRGB (Fig. 2) and YC$_b$C$_r$ (Fig. 5), in terms of uniformity; if anything, CIE L*u*v* and CIE L*a*b* CTFs are less uniform across the luminance range. An explanation is that CIE L*u*v* and CIE L*a*b* were largely optimized for color differences on the equiluminance plane, whereas our analysis focuses on CTF uniformity across luminance levels.

Both of these color spaces also require more bits than sRGB and YC$_b$C$_r$ to represent the same color gamut. This is because CIE L*u*v* and CIE L*a*b* were not designed for coding applications, thus not meant to optimize bits usage.
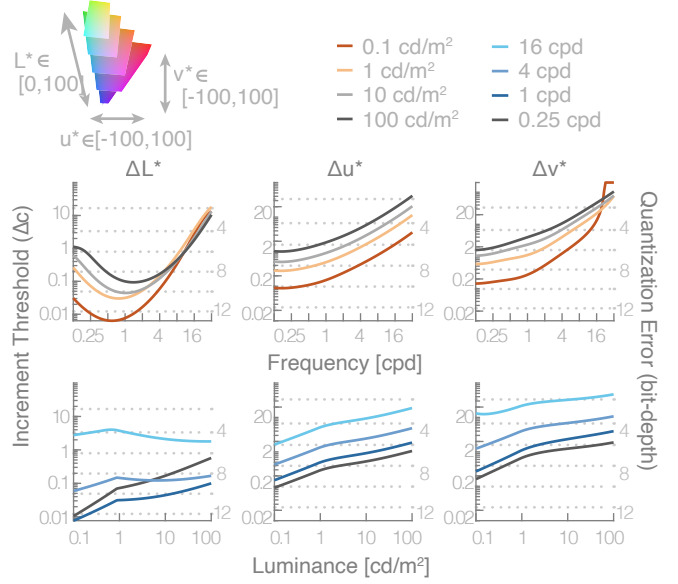
IC$_T$C$_P$

IC$_T$C$_P$ is an HDR color space designed to address the shortcomings of PQ-YC$_b$C$_r$ in representing saturated colors [9]. Color coordinates in IC$_T$C$_P$ are non-linear transformations of L, M, and S-cone responses, using PQ as the non-linearity. The channels are luma (I), blue-yellow (C$_T$) and red-green (C$_P$). I channel values are in the range [0,1] while C$_T$ and C$_P$ are in [-0.5,0.5]. The thresholds for IC$_T$C$_P$ (Fig. 9) have the same shape as those for PQ-YC$_b$C$_r$ (Fig. 6), but the thresholds for chroma channels are considerably higher. Consistent with the recommendation in [9], C$_T$ and C$_P$ require fewer bits compared to I channel. Based on our results, specifically, C$_T$ requires 8 bits and C$_P$ requires 9 bits to stay below detection threshold, whereas C$_b$ and C$_r$ require up to 11. Chroma sub-sampling is possible with IC$_T$C$_P$, since the increment thresholds for C$_T$ and C$_P$ are much higher at high spatial frequencies. This means that IC$_T$C$_P$ could be more efficient than PQ-YC$_b$C$_r$ in encoding chroma information.

CTF Uniformity

For coding applications, it is useful to have a metric that characterizes the uniformity of CTFs across luminance levels. Here, we propose uniformity error ε,

$$\varepsilon_i(f) = \log_2\left(\max_\ell \Delta c_i(f,\ell)\right) - \log_2\left(\min_\ell \Delta c_i(f,\ell)\right), \quad (4)$$

where $i \in \{1,2,3\}$ is the index of color channel. $\Delta c_i(f,l)$ is the increment threshold for i-th color channel at frequency f and luminance ℓ. The maximum and minimum thresholds are taken over all luminances ℓ within the color gamut. The units of the reported error are bits.

Table 1 shows CTF uniformity errors for 1 cpd, which is close to the peak sensitivity at most luminance levels, and for 4 cpd, which represents higher frequencies. The uniformity errors confirm that, indeed, CTF uniformity varies with spatial frequency, and that small uniformity error at one spatial frequency
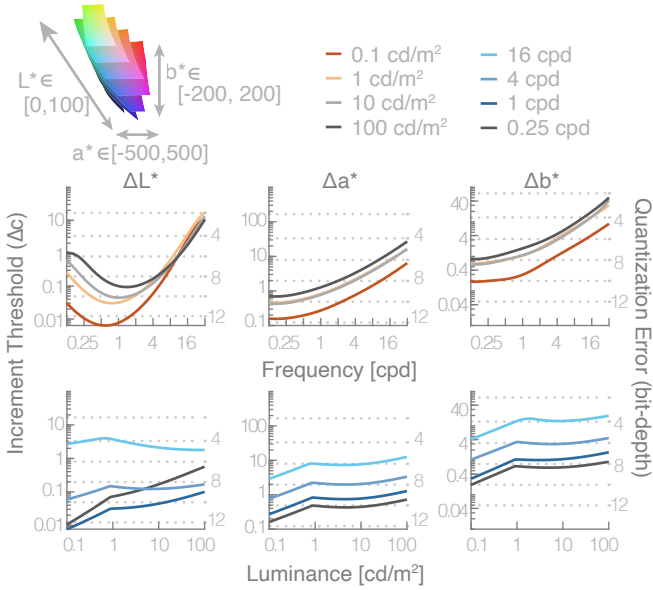
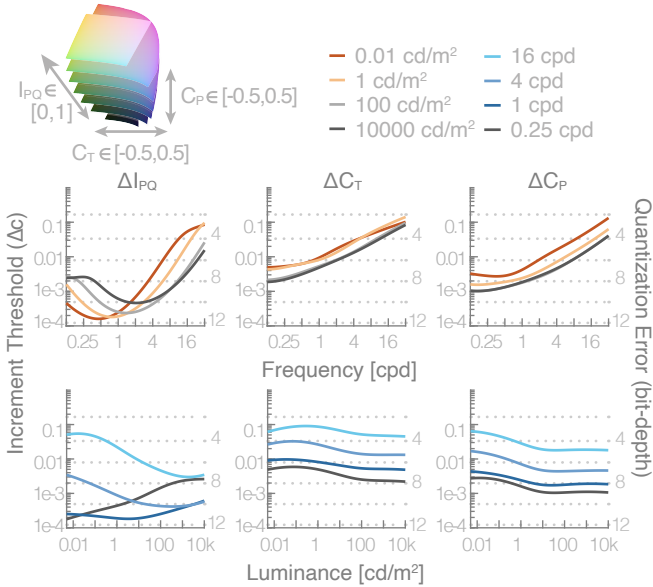Figure 8: CIE L*a*b* using ITU-R BT.709 color primaries [22].

Figure 9: IC_TC_P using ITU-R BT.2020 color primaries [10].

does not guarantee small error at a different spatial frequency.

The CTF uniformity errors are very large for linear RGB color spaces. This is why computing color differences in linear color spaces should be avoided [25]. The errors across other color spaces are comparable. For all color spaces, the achromatic channel is the least uniform. For PQ-RGB, the blue channel is more uniform than red and green. YC_bC_r is quite uniform, outperforming CIE L*a*b* and CIE L*u*v* for SDR, as well as IC_TC_P for HDR. C_P of IC_TC_P is less uniform than C_T, implying that C_P should be treated with more care than C_T for chroma subsampling or required bit-depth.

		1 cpd	4 cpd
HDR	Linear RGB	18.7, 19.3, 18.4	15.7, 15.2, 16.3
	PQ-RGB	0.97, 1.53, 0.60	2.42, 3.09, 1.86
	PQ- YC_bC_r	1.72, 0.80, 1.67	3.38, 1.17, 1.55
	IC_TC_P	1.74, 0.99, 1.42	3.40, 1.30, 2.01
SDR	Linear RGB	8.15, 8.24, 8.07	6.12, 6.06, 6.15
	sRGB	3.27, 3.36, 3.18	1.23, 1.18, 1.25
	YC_bC_r	3.36, 2.22, 1.90	1.16, 1.59, 1.81
	CIE L*u*v*	3.71, 3.63, 3.88	1.50, 3.51, 3.12
	CIE L*a*b*	3.71, 2.23, 2.54	1.50, 2.13, 2.09

Table 1: Uniformity errors in commonly used color spaces. Different color spaces order the chroma channels differently, and care should be taken when comparing across color spaces (e.g., C_r of YC_bC_r represents the red difference and is analogous to u*, not v*, of CIE L*u*v*).

General Discussion

We presented CTFs in SDR (sRGB, YC_bC_r, CIE L*u*v*, CIE L*a*b*) and HDR (PQ-RGB, PQ-YC_bC_r, IC_TC_P) color spaces based on the spatio-chromatic CSF in [14] and determined the minimum bit-depth required to keep quantization errors below visible thresholds. For all color spaces, we found that bits should be unevenly distributed across color channels in order to ensure quantization artifacts stay below threshold (e.g., for IC_TC_P, 12 bits for I, 8 bits for C_T, 9 bits for C_P). It should be noted that the CTFs were computed for detecting color differences from D65, meaning that the recommendations may not hold for differences from other colors, especially highly saturated ones.

We also proposed a metric for assessing the uniformity of CTFs. For both HDR and SDR, we found that YC_bC_r CTFs tends to be the most uniform. CIE L*a*b* and CIE L*u*v* CTFs were not very uniform, likely because they were based on color differences on the equiluminance plane (MacAdam ellipses) while our error metric focuses on the CTFs uniformity across the luminance range. Indeed, our analysis is limited to CTF uniformity across luminance levels. It neither accounts for perceptual uniformity across chromaticities, nor for uniformity of CTFs on other background colors. A more comprehensive analysis with a global uniformity metric would be an informative line of future research.

A remaining question is how to generalize the recommendation on required bit-depth across spatial frequencies. Here, we provided CTFs as a function of both spatial frequency and luminance. For real-life applications, where the image consists of multiple spatial frequencies, the CTFs must be combined to yield a single threshold per luminance. One possibility is to compute the minimum-threshold envelope over the spatial frequencies; this is the same conservative assumption made for PQ and other transfer functions, where the sensitivity is taken to be the peak sensitivity at each spatial frequency [5, 6]. This assumption, however, could be too conservative for many applications. If a specific application is well-understood, CTFs could be combined using a less conservative method that is customized to the problem. For example, if the main concern is the visibility of banding artifacts, it is possible to combine the spatial frequencies according to the Fourier spectrum of banding artifacts [26, 27].

In summary, our work highlights the shortcomings of existing commonly used color spaces, and provides a foundation for deriving more efficient perceptual transfer functions that can be

used in digital imaging applications.

Acknowledgments

References

[1] F. W. Campbell and J. G. Robson. Application of Fourier analysis to the visibility of gratings. *J. Physiol.*, 197(3):551, 1968.

[2] F. L. Van Nes and M. A. Bouman. Spatial modulation transfer in the human eye. *J. Opt. Soc. Am.*, 57(3):401–406, March 1967.

[3] K. T. Mullen. The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *J. Physiol.*, 359:381–400, February 1985.

[4] J. Rovamo, O. Luntinen, and R. Näsänen. Modelling the dependence of contrast sensitivity on grating area and spatial frequency. *Vision Res.*, 33(18):2773–88, December 1993.

[5] Rafał Konrad Mantiuk, K. Myszkowski, and H. P. Seidel. Lossy compression of high dynamic range images and video. In *Human Vision and Electronic Imaging XI*, volume 6057, pages 311–320. International Society for Optics and Photonics, SPIE, 2006.

[6] Scott Miller, Mahdi Nezamabadi, and Scott Daly. Perceptual signal coding for more efficient usage of bit codes. *SMPTE Motion Imaging Journal*, 122(4):52–59, 2013.

[7] P. G. J. Barten. *Contrast sensitivity of the human eye and its effects on image quality*. SPIE Press, 1999.

[8] BT Series. Worldwide unified colorimetry and related characteristics of future television and imaging systems. 1998.

[9] T. Lu, F. Pu, P. Yin, T. Chen, W. Husak, J. Pytlarz, R. Atkins, J. Frhlich, and G. M. Su. ITP colour space and its compression performance for high dynamic range and wide colour gamut video distribution. *ZTE Communications*, 14(1):32–38, 2019.

[10] BT Series. Parameter values for ultra-high definition television systems for production and international programme exchange. 2015.

[11] S. J. Anderson, K. T. Mullen, and R. F. Hess. Human peripheral spatial resolution for achromatic and chromatic stimuli: limits imposed by optical and retinal factors. *J. Physiol.*, 442(1):47–64, 1991.

[12] Y. J. Kim, A. Reynaud, R. F. Hess, and K. T. Mullen. A normative data set for the clinical assessment of achromatic and chromatic contrast sensitivity using a qCSF approach. *Investig. Ophthalmol. Vis. Sci.*, 58(9):3628–3636, 2017.

[13] A. M. Derrington, J. Krauskopf, and P. Lennie. Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.*, 357(1):241–265, 1984.

[14] Rafał Konrad Mantiuk, Minjung Kim, Maliha Ashraf, Qiang Xu, M. Ronnier Luo, Jasna Martinovic, and Sophie Wuerger. Practical color contrast sensitivity functions for luminance levels up to 10'000 cd/m^2. In *28th Color Imaging Conference*. IS & T, November 2020.

[15] S. M. Wuerger, M. Ashraf, M. Kim, J. Martinovic, M. Pérez-Ortiz, and R. K. Mantiuk. Spatio-chromatic contrast sensitivity under mesopic and photopic light levels. *J. Vision*, 2020.

[16] K. J. Kim, R. K. Mantiuk, and K. H. Lee. Measurements of achromatic and chromatic contrast sensitivity functions for an extended range of adaptation luminance. In *Human Vision and Electronic Imaging XVIII*, volume 8651, pages 319–332. International Society for Optics and Photonics, SPIE, 2013.

[17] M. Kim, M. Ashraf, M. Pérez-Ortiz, J. Martinovic, S. M. Wuerger, and R. K. Mantiuk. Contrast sensitivity functions for HDR displays. In *London Imaging Meeting*, London, UK, April 2020.

[18] R. K. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM T. Graphic*, 30(4):40:1–40:14, July 2011.

[19] S. M. Wuerger, A. B. Watson, and A. J. Ahumada, Jr. Towards a spatio-chromatic standard observer for detection. In *P. Soc. Photo-Opt. Ins.*, volume 4662, pages 159–172, June 2002.

[20] Q. Xu, M. R. Luo, and D. Sekulovski. Investigation of Spatial Chromatic Contrast around 5 Colour Centres. In *London Imaging Meeting*, 2020.

[21] CIE170-1:2006. Fundamental chromaticity diagram with psychological axes - part 1. Technical report, Central Bureau of the Commission Internationale de l'Éclairage, 2006.

[22] BT Series. Parameter values for the HDTV standards for production and international programme exchange. 2015.

[23] Charles Poynton. *Digital Video and HD: Algorithms and Interfaces*. Elsevier, 2012.

[24] M. T. Pourazad, C. Doutre, M. Azimi, and P. Nasiopoulos. HEVC: The new gold standard for video compression: How does hevc compare with h. 264/avc? *IEEE Consumer Electronics Magazine*, 1(3):36–46, 2012.

[25] R. K. Mantiuk. Practicalities of predicting quality of high dynamic range images and video. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 904–908. IEEE, September 2016.

[26] G. Denes, G. Ash, and R. K. Mantiuk. A visual model for predicting chromatic banding artifacts. In *Proc. SPIE*, 2019.

[27] Minjung Kim, Maryam Azimi, and Rafał Konrad Mantiuk. Perceptually motivated model for predicting banding artefacts in high-dynamic range images. In *28th Color Imaging Conference*. IS & T, November 2020.

Author Biography

Minjung Kim is a postdoctoral research associate at the Department of Computer Science and Technology and Robinson College, University of Cambridge. Her PhD is from New York University and York University, Canada. Her research interests include the perception of color, light, and shape, and the computational modeling of visual perception. `https://www.minjung.ca`

Maryam Azimi is a postdoctoral research fellow of the Natural Sciences and Engineering Research Council of Canada with the Department of Computer Science and Technology, University of Cambridge. She received her MASc and PhD degrees in electrical and computer engineering from the University of British Columbia, Vancouver, Canada, in 2014 and 2019, respectively. She is an active member of the Standard Council of Canada, MPEG and the IEEE signal processing society.

Rafał K. Mantiuk is Reader (Associate Professor) at the Department of Computer Science and Technology, University of Cambridge (UK). He received PhD from the Max-Planck-Institute for Computer Science (Germany). His recent interests focus on computational displays, novel display technologies, rendering and imaging algorithms that adapt to human visual performance and viewing conditions in order to deliver the best images given limited resources, such as computation time, bandwidth or dynamic range. He contributed to early work on high dynamic range imaging, including quality metrics (HDR-VDP), video compression and tone-mapping. Further details: `http://www.cl.cam.ac.uk/~rkm38/`.