59.2: Defining Dynamic Range

Greg Ward

Dolby Canada, Vancouver, BC, Canada

Abstract

We introduce a new metric for dynamic range of displays that closely corresponds with human perception in practical settings. The Number of Distinguishable Gray*s is a count of the visible luminance steps from the deepest black to the brightest white a display produces under known ambient conditions.*

1 Introduction

Progress in high dynamic range displays relies on adequate methods for comparing and understanding differences between competing technologies and products. This is currently hampered by the lack of a clear definition of dynamic range, usually represented as the maximum contrast ratio or CR. The Video Electronics Standards Assoc. (VESA) 2.0 standard [1,2] specifies measurement conditions for full-screen white and full-screen black. The ratio of these two values is most often used by manufacturers to denote maximum contrast for their displays, despite complaints by both consumers and manufacturers that the specification permits excessive tampering with display settings prior to such measurements to obtain unrealistic CR numbers [3].

A more critical problem is the precise meaning of "black" in such measurements, which ignores the effect of ambient lighting and the importance of visible quantization in determining whether the measured minimum is useful in any practical context. Emerging high dynamic range displays may have 10, 12, and even 16-bit/channel inputs, and a metric that quantifies the effect of bit depth and gamma response on perception is essential. We address each of these issues in our recommended approach, and suggest a new dynamic range metric that incorporates ambient lighting and quantization, which we call the *Number of Distinguishable Grays* (NDG).

The question NDG answers is, "How many gray levels will a standard observer be able to distinguish between the darkest black and the brightest white in a given ambient environment?" We submit that this is an easily understood concept from a buyer's perspective and a meaningful metric from a vendor's standpoint, which defines dynamic range more reliably than the contrast ratio. Furthermore, NDG values are perceptually well-spaced by design; i.e., a display with twice the NDG value has roughly twice the apparent dynamic range.

2 NDG Method

Determining the Number of Distinguishable Grays for a display requires the following three pieces of information, comprised of two measurements and a specified condition:

1. The gray level luminance response of the display for identical RGB inputs, $L_d(i)$, including the range of i (e.g., 0-1023 for 10-bit/channel input).
2. The average hemispherical reflectance of the display surface.
3. The presumed ambient lighting for this NDG evaluation.

If the display is assumed to be in a completely darkened environment, the hemispherical reflectance is not needed. However, we strongly recommend that any new standard require at least one NDG evaluation under typical indoor lighting conditions, either home lighting (around 50 lux vertical illuminance) for television units, or office lighting (around 200 lux) for video display terminals. This may be given in addition to the "best case" darkened viewing environment. Even for televisions, the worst case for front reflections may be dozens of times greater than the proposed 50 lux level.

The NDG value may then be computed with the following formula:

$$NDG = \sum_{i=1}^{i_{\max}} \min\left(\frac{L_d(i) - L_d(i-1)}{D(L_d(i) + L_a)}, 1 \right)$$

where L_a is the reflected display luminance due to ambient lighting, computed as the vertical ambient illuminance times the hemispherical front reflectance, divided by π. The *D(L)* function is an agreed-upon *threshold versus intensity* or *t.v.i.* curve based on psychophysical data. A few such curves have been proposed, such as Barten's formula employed in the DICOM standard [4], Ferwerda et al.'s perceptually-based tone-mapping [5], and Lubin and Pica's perceptual quantizer [6]. The choice of t.v.i. curve is not critical, so long as it is well-specified and adhered to by the standard.

The *min(f,1)* function in the NDG formula serves an important purpose, which is to count fractional steps for sub-threshold jumps, while never giving more than one NDG per control quanta. In other words, we will not give an NDG of 4096 to a 12-bit display unless an observer can see every one of those steps. Nor will we give an NDG value of 256 to an 8-bit display if the quanta near white take supra-threshold steps while the bottom 20 quanta are lost in unwanted screen reflections.

Display	Screen Reflectance	Maximum Luminance	Ambient Level (incident illuminance)	Contrast Ratio	Number of Distinguishable Grays
17” LCD **A** 8-bit input	2%	200 cd/m^2	0 50 lux 200	400 244 113	240 237 234
17” LCD **B** 8-bit input	1%	200 cd/m^2	0 50 lux 200	400 303 176	240 238 236
17” LCD **C** 8-bit input	1%	400 cd/m^2	0 50 lux 200	400 345 244	242 242 241
17” LCD **D** 10-bit input	1%	400 cd/m^2	0 50 lux 200	400 345 244	351 348 338

Table 1. The CRs and NDGs of four hypothetical displays under three ambient conditions.

In a simple illustration of the NDG metric, we compared four hypothetical displays, measured at different ambient lighting levels. We assume an sRGB response from each, and a VESA CR of 400. LCD **A** has a screen reflectance of 2% and LCD **B** has a screen reflectance of 1%. Both **A** and **B** have the same maximum luminance of 200 cd/m^2. In addition to having a lower screen reflectance, LCD **C** has a boosted maximum luminance of 400 cd/m^2, and LCD **D** provides a 10-bit input. (These specifications are chosen for illustrative purposes rather than to match any actual display.) Our *D(L)* curve was adapted from [5], as described in the following subsection. The “results” of our virtual test are summarized in Table 1.

As we can see from our table, the NDG values for displays A, B, and C are limited primarily by their 8-bit inputs, whose delta between levels is greater than the visible threshold over most of the range. Ambient lighting and display reflectance have only a minor influence by comparison, and it is only when we move to the 10-bit input of display D that the NDG becomes an actual count of visible thresholds over the display luminance range. It is important that the 10-bit response follows the sRGB curve. If we had switched to a linear response when we increased the number of bits to 10, we would have only have achieved a NDG of 298 at L_a=0. This is because the linear step sizes would exceed the $D(L_a)$ value at the bottom end of the range.

2.1 Threshold vs. Intensity

As noted earlier, calculating NDG from display output requires a human threshold vs. intensity curve relating threshold response to adaptation luminance. Psychophysical studies have established this non-linear relationship, shown for rods and cones in Figure 1 [5]. The piecewise fit to the combined curve is written in Table 2 [7].

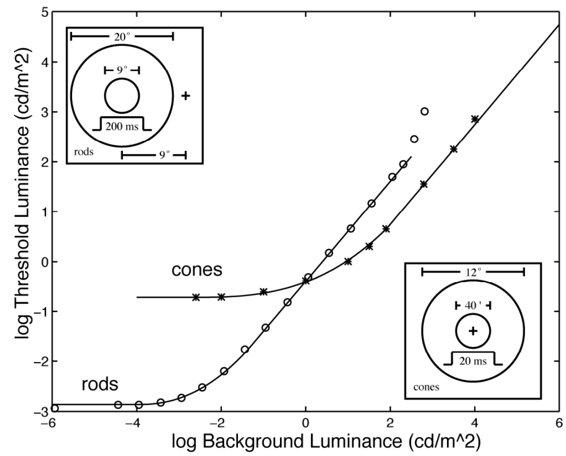

Figure 1. Plot of threshold vs. intensity (t.v.i.) for human vision.

log10 of Just Noticeable Difference tvi(L_a)	**Applicable Adapted Luminance Range**
-2.86	$\log_{10}(L_a) < -3.94$
$(0.405 \log_{10}(L_a) + 1.6)^{2.18} - 2.86$	$-3.94 \le \log_{10}(L_a) < -1.44$
$\log_{10}(L_a) - 0.395$	$-1.44 \le \log_{10}(L_a) < -0.0184$
$(0.249 \log_{10}(L_a) + 0.65)^{2.7} - 0.72$	$-0.0184 \le \log_{10}(L_a) < 1.9$
$\log_{10}(L_a) - 1.255$	$\log_{10}(L_a) \ge 1.9$

Table 2. Numerical fit to t.v.i. shown in Figure 1.

There is general agreement that Ferwerda's thresholds are on the high side, and banding may be visible in gradient images at thresholds as low as 1/10th of the reported t.v.i. curve (corresponding to 0.6% steps in the photopic region). The probable reason for this discrepancy is that the experiments from which this function is derived use a pulsing target on a constant background, and the thresholds for such transient stimuli are higher than those for static stimuli. We therefore recommend subtracting 0.95 from the formulae on the left side of Table 2, which is equivalent to dividing the JND threshold luminances by a factor of 9. This brings our t.v.i. function into better agreement with the Barten model recommended for radiology by the Digital Imaging and Communications in Medicine (DICOM) standards body, as shown in Figure 2. We did not follow the DICOM recommendation directly because their fit to Barten's data only covers part of the visible luminance range. This range, 0.05:4000 cd/m^2, while greater than any conventional display, is on par with dual modulator displays just entering the market, and may be exceeded in the near future. It is therefore preferable to either extend the Barten fit to cover the full range of human vision, or adjust Ferwerda's model to match a more realistic threshold. We follow the latter strategy, although either approach will work.

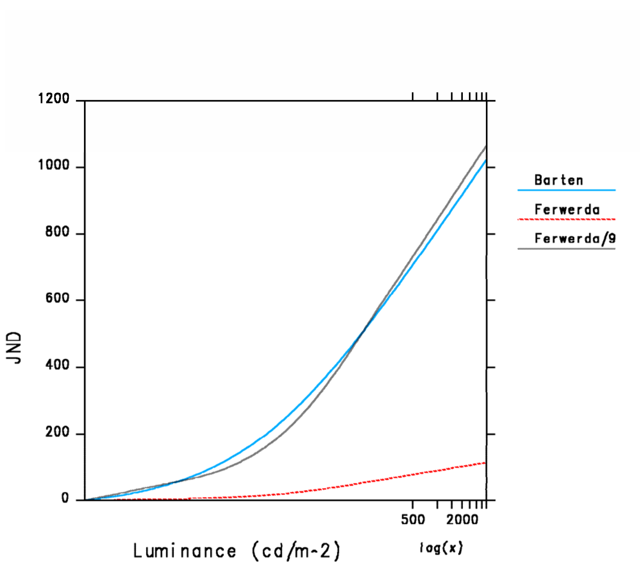

Figure 2. The DICOM fit to Barten's model, and our adjustment of Ferwerda's.

We use this modified Ferwerda model for *D(L)*:

$$D(L) = 10^{tvi(L) - 0.95}$$

Using a threshold model places an upper limit on the acceptable quantization error, and offers a staircase scale corresponding to visible luminance differences over a given display's dynamic range independent of quantization. This allows us to distinguish between a display that covers a wide dynamic range, but in a region where humans can barely see (e.g., below 10^{-2} cd/m^2) versus a display that covers a range where we see well.

3 Conclusion

Knowledgeable consumers and many manufacturers agree that the current practice for measuring maximum contrast ratio, if not the VESA 2.0 standard itself, is suspect to the point of being meaningless for comparisons. One solution is to shore up the VESA standard by requiring that manufacturers fix their monitor settings to their defaults prior to measurements. Product specifications should also include the effects of ambient illumination on delivered contrast, perhaps offering a few different levels so the consumer may judge appropriately for their intended application.

However, fixing the contrast ratio measurement alone does not provide for intelligent evaluation of displays with greater bit depth and controllable dynamic range. For emerging high dynamic range displays, we need a sensible way to quantify the delivered range of visible luminances. In this paper, we have suggested the *Number of Distinguishable Grays* specification. By measuring the gray output levels of a display for each RGB=(i,i,i) input and applying the given formula, one arrives at a quantity with good correspondence to the actual visible differences, accounting simultaneously for the effects of brightness, contrast, ambient lighting, and quantization. With a modest education effort, we believe the NDG concept is intuitive and will be appreciated by consumers and professionals alike.

4 References

[1] "Flat Panel Display Measurements" 2.0 avail. from www.vesa.org

[2] "Flat Panel Display Measurements Update" avail. from www.vesa.org

[3] Tom Mainelli, "Tested contrast ratios rarely conform to vendors' specs," PC World e-zine www.pcworld.com/news/article/0,aid,110483,pg,1,00.asp, May 2003.

[4] "Part 14: Grayscale Standard Display Function," *Digital Imaging and Communications in Medicine (DICOM)*, National Electrical Manufacturer's Assoc.

[5] James Ferwerda, S. Pattanaik, P. Shirley and D. Greenberg. "A Model of Visual Adaptation for Realistic Image Synthesis," *Proceedings of ACM SIGGRAPH '96*

[6] Jeffrey Lubin, A. Pica, "A Non-Uniform Quantizer Matched to Human Visual Performance," *Soc. for Information Display Digest*, 1991.

[7] Greg Ward Larson, H. Rushmeier, C. Piatko, "A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes," *IEEE Transactions on Visualization and Computer Graphics*, Vol. 3, No. 4, December 1997. radsite.lbl.gov/radiance/papers/lbnl39882/tonemap.pdf