Measuring the impact of flare light on Dynamic Range

Norman Koren; Imatest LLC; Boulder, CO USA

Abstract

The dynamic range (DR; defined as the range of exposure between saturation and 0 dB SNR) of recent High Dynamic Range (HDR) image sensors, can be extremely high: 120 dB or more. But the dynamic range of real imaging systems that include lenses is limited by veiling glare (susceptibility to flare light from reflections inside the lens), and hence rarely approaches this level. Standard veiling glare measurements, such as ISO 18844, made from charts with black cavities on white fields, yield numbers (expressed as a percentage of the pixel level in nearby light areas) that are much worse than expected for actual camera dynamic range.

Camera dynamic range is typically measured from grayscale charts, and is strongly affected by veiling glare, which is a function of the lens, chart design, and the surrounding field. Many HDR systems employ tone mapping— which enables HDR scenes to be rendered in displays with limited dynamic range by compressing (flattening) tones over large areas while attempting to maintain local contrast in small areas. Measurements of tone-mapped images from standard grayscale charts often show low contrast over a wide tonal range, and give no indication of local contrast, which is especially important for the automotive and security industries, where lighting is uncontrolled and the visibility of low contrast features in shadow regions is critical.

We discuss the interaction between veiling glare and dynamic range measurements and we propose a new transmissive test chart and dynamic range definition that directly indicates the visibility of low contrast features over a wide range of scene brightness.

Introduction

Although several authors have discussed the relationship between veiling glare and dynamic range [1, 2], they are usually tested separately. This may be reasonable for image sensors with limited dynamic range, but is no longer tenable for the new generation of high dynamic range (HDR) sensors, where camera performance is limited by veiling glare from the lens.

A dynamic range race (somewhat akin to the automotive horsepower race of the 1950s) seems to be developing where marketing departments expect engineers to produce camera dynamic range measurements comparable to HDR sensor specifications. Unfortunately this can lead to highly misleading marketing material.

Veiling glare measurements

Veiling glare is a measure of the average susceptibility to flare light—fogging or ghost images caused by light originating inside or near the image. Flare can be thought of as having three components: 1. overall image fogging, 2. ghost images, typically arising from small, intense light sources (important, but beyond the scope of this paper), 3. stray light that decays with distance from a point source, characterized by a glare spread function (GSF). The GSF approximates an exponential decay that may vary over the image surface. As we shall see, the GSF can cause errors in dynamic range measurements.

In ISO 18844:2017, measurement type C [3], veiling glare is measured from charts that contain black cavities in a larger white field.

$$Veiling\ glare\ (\%) = 100\% \times Y_B/Y_W \quad \text{, where} \tag{1}$$

Y_B is the average luminance inside the dark cavity and Y_W is the average luminance of specified white areas near the cavity. Both Y-values are derived from the linear or linearized pixel level, best converted from raw format with gamma = 1.

Although it is normally reported as a percentage value, it can also be expressed in decibels (dB) for convenient comparison with dynamic range.

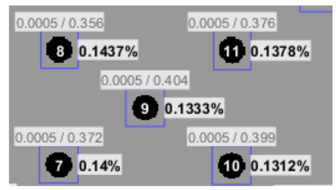

Figure 1. Portion of ISO 18844 veiling glare density

$$Veiling\ glare\ (dB) = 20\ log_{10}(Y_W/Y_B) \tag{2}$$

This measurement ignores effects of the GSF. For a high quality camera and lens (Sony A7Rii with the 5-element Canon 90mm T/S-E f/2.8 lens at f/5.6) the mean veiling glare is 0.187% = 55 dB. With a consumer-grade zoom lens (Canon 75-300 f/4-5.6 at f/5.6) veiling glare was 0.130% = 58 dB.

Veiling glare measurements can be thought of as an extreme worst-case of dynamic range. Measured values are far below sensor dynamic range. This is hardly a realistic use case, but as we shall see, neither is the standard technique for measuring image sensor dynamic range.

Sensor and camera dynamic range

Image sensor dynamic range, defined at the ratio of illumination just under sensor saturation to illumination where SNR = 1 (0 dB), is typically measured from a set of flat field images for calibrated light levels projected on the image sensor, using equation (18) from section 2.4 of the EMVA 1288 standard [4]. Each image used for this measurement has zero dynamic range. This is a best case measurement since there are no losses from lens flare.

Camera dynamic range is typically measured from images of transmissive grayscale test charts that have a relatively large number of patches (20-36) with a wide range of densities ($D_{max} - D_{min} \geq 3$). The patches may be arranged in a linear or circular pattern. Although circular patterns are preferred because they are less susceptible to lens vignetting and are easier to frame (especially with distorted images), linear charts are quite common.

For linear sensors, image sensor dynamic range can also be measured from transmissive grayscale charts if the image is minimally processed (no color correction, noise reduction, gamma curve, offsets, etc.). The image is analyzed and noise is fit to an equation derived from section 2.2 of the EMVA 1288 standard.

$$\sigma = noise = \sqrt{dark\ noise^2 + k * pixel\ level} \tag{3}$$

Camera dynamic range is defined by two criteria: the range of illumination where (1) the scene-referenced Signal-to-Noise Ratio $SNR_{scene\text{-}ref}$ is above a specified amount (10 or 20dB for "high" quality; 1 or 0dB for "low" quality, which corresponds to sensor dynamic range measurements), and (2) the slope of the log pixel level vs. log exposure curve is greater than a specified fraction of the maximum slope.

$SNR_{scene\text{-}ref}$ is calculated by dividing the standard SNR by the slope of the tonal response curve, resulting in a value that would be visible in the scene. A factor of two change in illumination is called one f-stop (or zone or EV for Exposure value, equivalent to $\log_2(exposure)$) by photographers, and can be used as a perceptual unit for noise measurement.

$$f\text{–}stop\ noise = \frac{noise\ in\ pixel\ levels}{d(pixel\ level)/d(f\text{–}stop)}\ ,\quad \text{where} \tag{4}$$

$$d(f\text{-}stop) = d\left(\log_2(exposure)\right) = 1.443\, d\left(\log_e(exposure)\right) = 1.443\,\frac{d(exposure)}{exposure} \tag{5}$$

When we drop the 1.443 to maintain compatibility with older calculations we observe that the inverse of *f–stop noise* is, in fact, the scene-referenced SNR.

$$SNR_{scene\text{–}ref} = \frac{exposure}{(noise\ in\ pixel\ levels)}\frac{d(pixel\ level)}{d(exposure)} \tag{6}$$

Veiling glare degrades dynamic range by fogging dark areas of the image. Although it increases the (non-scene referenced) SNR measured in the camera, it reduces $SNR_{scene\text{-}ref}$ in the darker portions of the image by decreasing *d(pixel level)/d(exposure)*.

The second dynamic range criterion, the range of illumination where the slope *d(pixel level)/d(exposure)* is greater than a specified fraction of the maximum (we have been using 0.075, which may be lower than optimum), is called the slope-based dynamic range, DR_{slope}. It is frequently larger than dynamic range based on $SNR_{scene\text{-}ref}$, i.e., $SNR_{scene\text{-}ref}$ may be well under 1 (0 dB) at the lower limit of DR_{slope}, meaning no features of interest are likely to be visible.

Typical results

Figure 2 shows the density response and scene-referenced SNR for the Canon EOS 6D— a full-frame Digital SLR with a high quality (linear, non-HDR) sensor— at ISO 100 with a 100mm f/2.8 macro lens at f/8. It was converted to a 48-bit TIFF file with dcraw (which applies no noise reduction).

The upper plot shows log Pixel level as a function of exposure (-chart density) in units of decibels (dB = 20 x Optical Density). Gamma is close to the expected value for Adobe RGB (1/2.2 = 0.4545).

DR_{slope} is 82.9dB. At its lower limit $SNR_{scene\text{-}ref}$ is around −10dB for this image— much too low for image detail to be visible or results to be repeatable. Some engineers report this number as the "total" dynamic range of the system, which is true only in the crudest sense since no details are distinguishable where SNR is below 0dB. This practice is apparently encouraged by marketing departments, who want to report the highest possible dynamic range for their products. We discourage it.

The lower plot shows $SNR_{scene\text{-}ref}$ in dB. Dynamic range is 41.8 dB at high quality (20 dB), increasing to 66.3 dB for low quality (0 dB). The image sensor dynamic range, derived from a minimally processed version of this image, is 11.9 EV = 71.7 dB.

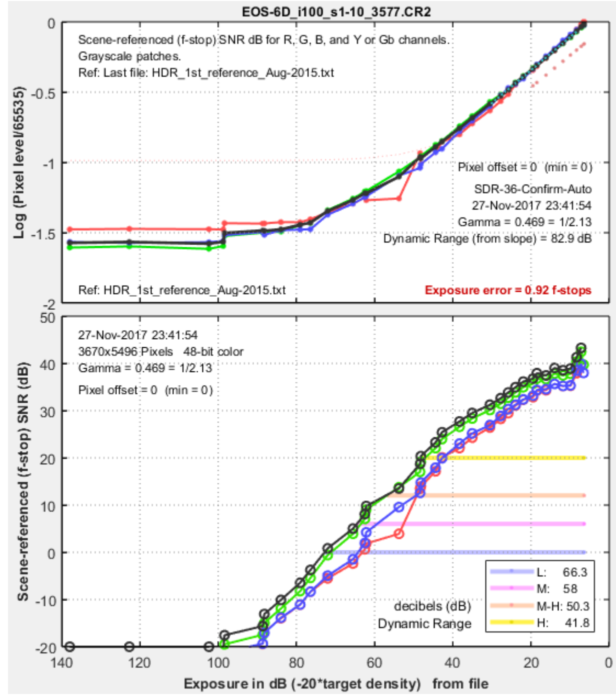

Figure 2. Typical results of dynamic range analysis

Problems with standard grayscale charts

Standard grayscale charts have an important issue that can affect measurement accuracy. Medium-range flare light that falls off with increasing distance from light patches, but may extend to large distances over the image, can affect dynamic range measurements. It is easiest to detect on linear charts, but it has a similar (and possibly worse) effect on circular charts.

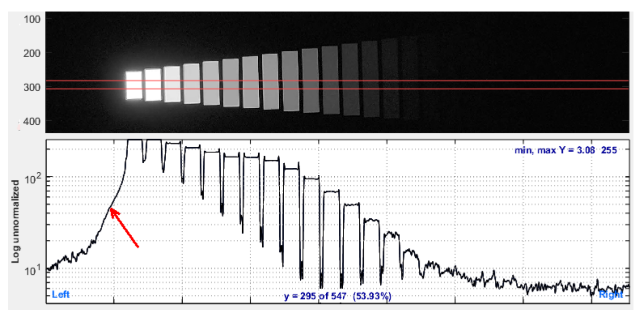

Figure 3. Xyla chart image and horizontal cross-section

The top of Figure 3 is an image of a DSC Labs Xyla chart, which has at least 21 precise patches with an optical density (OD) step of 0.3, from a security camera whose technical details we don't know. Below the chart image is a horizontal cross-section of the pixel levels, displayed logarithmically.

The brighter patches have flare light diffusing around them— very visible when the image is lightened, and quite obvious in the cross-section plot (see the red arrow). This flare light can affect readings in dark patches at large distances from the lightest patches. Note also that the patch levels seem to flatten out around x = 500 to 650. This is likely caused by tone mapping. There is no way of knowing the local contrast of small objects in this exposure range.

Figure 4 is a cross-section plot taken just below the chart (inside the parallel red horizontal lines), showing the effects of flare, which affects pixel levels at large distances from the brightest patches. As flare increases (for low quality lenses) it becomes likely that its effect at a distance will increase the slope-based dynamic range.

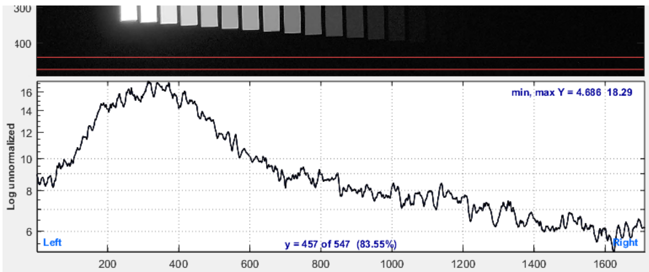

Figure 4. Cross section taken just below Xyla chart

Figure 5 shows the density response of this image on top and $SNR_{scene\text{-}ref}$ at the bottom. Dynamic range is quite good.

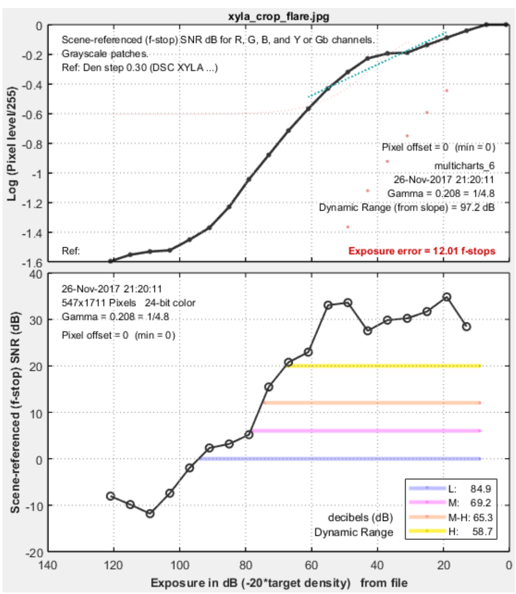

Figure 5. Density response and SNRscene-ref. for the Xyla image

The oddly-shaped density plot with low gamma (0.208, versus 0.454 for standard sRGB) indicates that tone-mapping has been applied. Tone mapping is intended to reduce *global* contrast while maintaining *local* contrast, but there is no way of determining how well local contrast is maintained with this type of test chart.

A new approach to dynamic range measurement

A practical measurement of dynamic range must answer the question, "How visible are low-contrast objects over a wide range of illumination?" A robust measurement should give a reliable indication of local contrast when tone-mapping has been applied and must be immune to errors caused by medium range flare light, as described above.

To meet these goals we have developed a new test chart called the ***Contrast Resolution chart*** that contains pairs of relatively low contrast (2:1 ratio = 6dB) gray patches inside twenty larger gray patches whose Optical Densities (ODs) range from base + 0.15 to base + 4.90 in steps of 0.25, equivalent to 95 dB. (If the lightest and darkest small patches were included, the total density range would be 5.05 OD = 101 dB.) Figure 6 illustrates the overall chart concept.

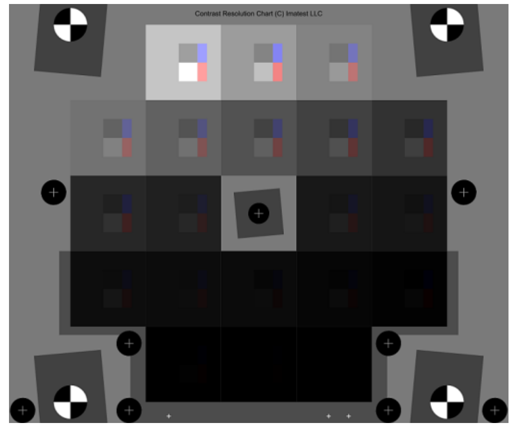

Figure 6: Contrast resolution chart concept

The physical Contrast Resolution chart is made from two layers of 8×10 inch color photographic film. The left side of the large patches are used for noise measurements because the small patches are usually too small for good noise statistics. The mean values of the pairs of light and dark grayscale patches (as well as the means of the blue and red patches, included for visual analysis) have the mean pixel level as the surrounding patch. This ensures that the small patches will have a minimal effect on tone mapping.

As with all transmissive dynamic range charts it must be photographed in a completely dark environment. Opaque black cloth (velvet or felt) should cover any objects that might reflect significant amounts of light back to the chart, which would compromise measurements.

Contrast resolution results

Figure 7 contains results for an image of the Contrast Resolution chart taken with the Sony A7Rii full-frame mirrorless camera at Exposure index = ISO 100. The lens was the 5-element Canon 90mm TS-E f/2.8 set to f/5.6. Sensor dynamic range, measured from a minimally-processed Contrast Resolution chart image, is 13 EV = 78 dB. Many of the results are identical to standard grayscale charts, but the results shown in magenta are uniquely defined for the Contrast Resolution chart.

The upper plot in Figure 7 contains two thick curves. The gray curve is the logarithm of the standard signal—the pixel level of the larger gray patches. The magenta curve (below the gray curve) is the logarithm of the **Contrast Resolution signal (S_{CR}),** defined as the difference in pixel level between the small light and dark gray patches (which have a 2:1 contrast ratio or 6 dB difference on the chart).

The lower plot shows the standard SNR as a gray curve and the **Contrast-Resolution SNR (SNR_{CR})**, defined as the Contrast-Resolution signal S_{CR} divided by the noise measured in the larger gray patch) as a bold magenta curve.

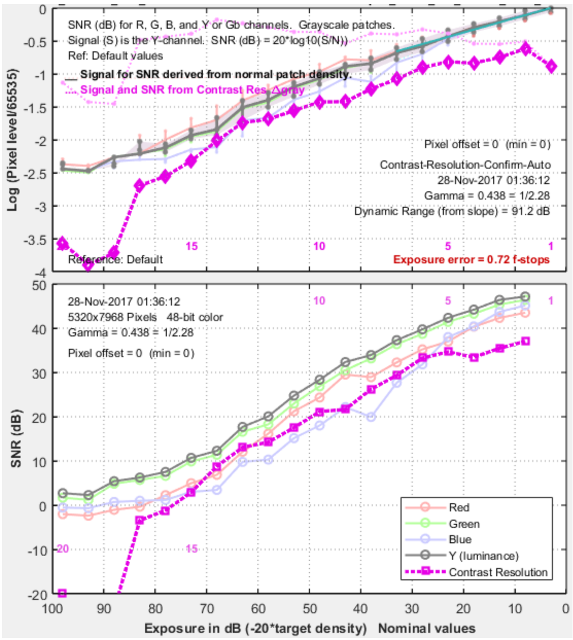

Figure 7. Results for the Contrast resolution for the Sony A7Rii, showing Contrast Resolution signal and SNR.

Figures 8 illustrates the visibility of low contrast features— the small light and dark gray patches. It is created by adjusting the mean pixel level of each large patch (including all inner patches) so that they all have identical mean pixel levels. (This would be meaningless with traditional grayscale charts.) We have had the best results using the Y (luminance) channel from CIE 1931 xyY space (Y is linear) to adjust the image.

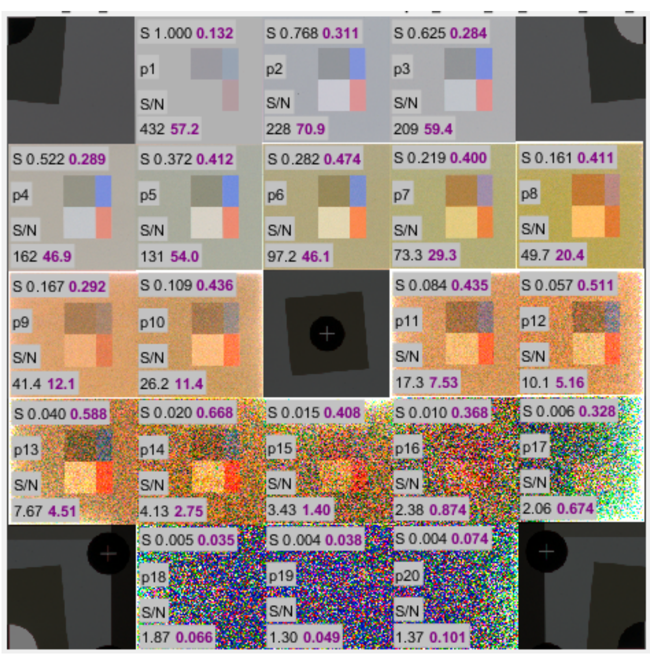

Figure 8. Image adjusted so large patches have equal average luminance level.

The first number in the top line of each patch is the standard signal level (normalized to the maximum level for the bit depth), displayed as the gray curve in the upper plot of Figure 7. The second number is the Contrast-Resolution (light-dark patch) signal level, S_{CR}, normalized to the standard signal level, displayed as a dotted magenta curve in Figure 7. This number is closely related to the visible patch contrast in Figure 8, which is effectively normalized.

The bottom line of each patch displays the standard SNR and Contrast-Resolution SNR, SNR_{CR}, (as ratios). These numbers are displayed in the lower plot of Figure 7 in dB as gray and magenta curves, respectively.

For this image, which maintains good contrast in the dark patches (as far as patch 17), feature visibility is dominated by SNR_{CR}. Patch 15, which is the darkest patch where the inner squares are clearly distinguishable has SNR_{CR} = 1.40 (3 dB).

When the chart was acquired with the Canon 75-300 f/4-5.6 zoom lens at f/5.6, the standard signal in the darker patches tended to be slightly lighter because of added veiling glare, but the Contrast-Resolution signal S_{CR} and visible color saturation were significantly reduced. SNR_{CR} is comparable to the 90mm lens image, but feature visibility is lower, especially in patches 13-15.

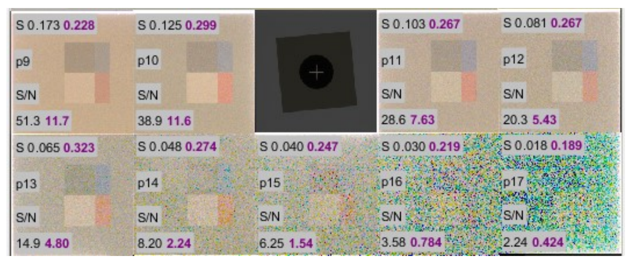

Figure 9. Portion of equal luminance image for the 75-300mm lens for patches 9-17.

For this image, the reduced Contrast-Resolution signal S_{CR} signal plays a significant role in reducing feature visibility. More perceptual testing is needed to develop a reliable visibility metric, but at least we have a clear path forward, which isn't available with standard grayscale charts that lack low contrast features.

Contrast Resolution with added veiling glare

Figure 10 shows the test chart photographed in front of a large LED lightbox. The light surrounding the chart (which we kept in its mount) significantly increases veiling glare.

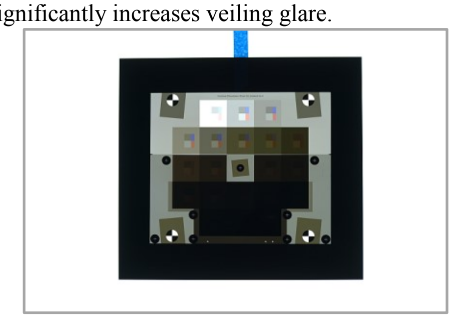

Figure 10. Contrast Resolution chart in front of large lightbox for measuring veiling glare.

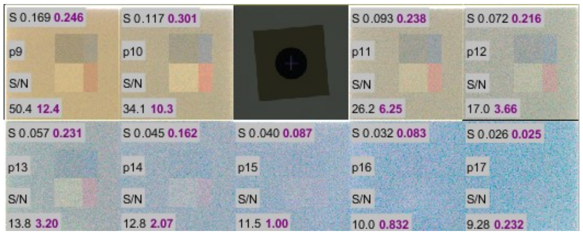

Figure 11. Portion of equal luminance image for the Contrast Resolution chart in front of large lightbox. 75-300mm lens, patches 9-17.

Note the decrease in visibility of low contrast features as we go from Figure 8 (a high quality prime lens) to Figure 9 (a consumer-grade zoom lens) to Figure 11 (the zoom lens with light surrounding the chart, shown in Figure 10). The extra flare light improves the standard SNR (the number on the lower-right of each patch, measured from the large gray area-only) starting in patch 14, but degrades the Contrast Resolution SNR. This is a good example of how using standard SNR measurements can be misleading in the presence of flare light. Unfortunately we found some problems with reflections when we tried this approach with camera phones that need to be close to the chart, so we can't generally recommend it.

HDR mode in a high quality camera phone

The Google Pixel phone has a switchable high dynamic range mode called HDR+ that performs some remarkable image processing [5] that includes noise reduction and tone mapping starting with multiple underexposed images that are resistant to highlight saturation. We used the Contrast Resolution chart to study the performance of HDR+.

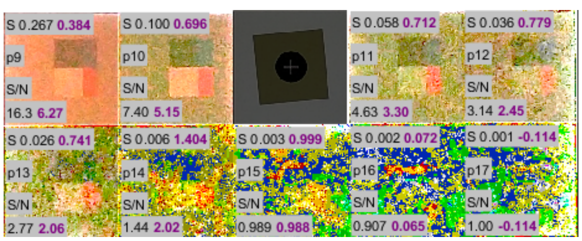

Figure 12. Portion of equal luminance image for the Google Pixel phone with HDR+ OFF (linear mode) for patches 9-17.

Figures 12 and 13 show crops of the equal luminance image for HDR+ OFF and ON, respectively. Differences are highly visible. HDR+ had much better visible appearance as well as SNR for both standard and Contrast-Resolution measurements. Note the blurring in patches 13 and 14 of the HDR+ image (Figure 13). Blurring (low-pass filtering) is a common strategy for reducing noise.

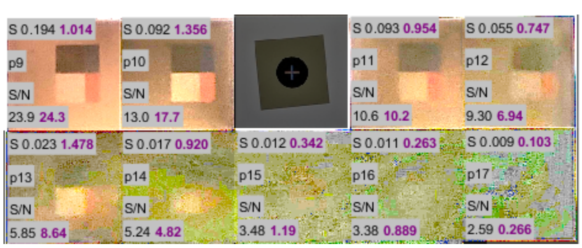

Figure 13. Portion of equal luminance image for the Google Pixel phone with HDR+ ON for patches 9-17.

Figures 14 and 15 show tonal response and SNR for HDR+ OFF and ON, respectively. Note that the upper (log pixel level) plots have different scales (minimum values of -4 and -3, respectively), so they differ by more than they appear at first glance.

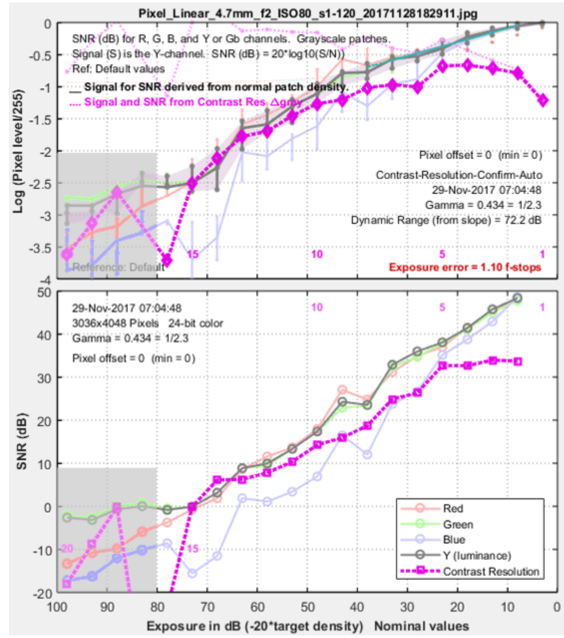

Figure 14. Contrast resolution results for the Google Pixel phone with HDR+ OFF (linear mode).

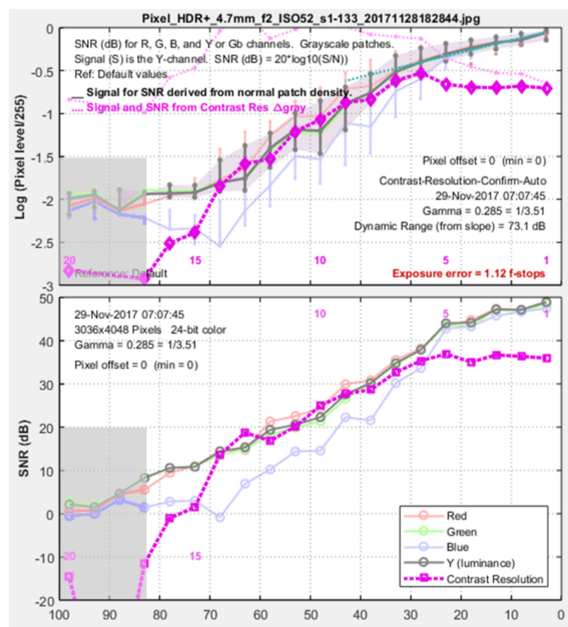

Figure 15. Contrast resolution results for the Google Pixel phone with HDR+ ON. Note that the scales of the upper plots of Figures 15 and 16 are different: minimum values = -4 and -3, respectively.

Portions of the two plots are grayed-out because either standard SNR or SNR_{CR} is well below zero, hence these results are neither reliable nor repeatable.

Gamma for HDR+ OFF is 0.434, close to the ideal encoding value for sRGB. Gamma for HDR+ ON is 0.285, indicative of much lower contrast (caused by tone mapping), but the Contrast Resolution signal is consistently higher by about 0.5 (log_{10} units), which shows that *local* contrast has increased despite the reduced gamma, which is a measurement of *global* contrast.

The lower plots in Figures 14 and 15 show the standard and Contrast Resolution SNR. Turning HDR+ ON improves the Contrast Resolution SNR by 7-10 dB.

Summary

Dynamic range needs to be defined as the range of exposure where the image contains useful detail. Any DR measurement that goes beyond this range is deceptive. As we have seen, standard dynamic range measurements don't always meet this standard. Little useful image information is visible at the usual lower SNR limit of 0 dB. Significant veiling glare or tone mapping can lead to erroneous results, which can have serious consequences for automotive imaging, where lives depend on good quality images.

A good test chart should be reasonably predictive of camera performance for a realistic range of use cases. We can't expect it to cover *all* cases.

Since low contrast objects are integral to an improved dynamic range definition, they should be a part of a test chart design. In the Contrast Resolution chart we use small light and dark squares with 2:1 contrast (6 dB difference) for this purpose. We can observe the visibility of the patches and correlate our observations with the Contrast Resolution signal (the difference between the two patches) and SNR_{CR}.

We have found many cases where the threshold of visibility corresponds to SNR_{CR} around 3-10 dB, but we have also found cases where visibility is reduced when the Contrast Resolution signal S_{CR} has been reduced, even though SNR_{CR} would have been adequate for higher S_{CR}. More perceptual work needs to be done to define the threshold of human visibility.

In traditional measurements SNR = 0 has been enshrined as the lower limit for dynamic range, even though little work has been done to correlate it with feature visibility. Our experience suggests that SNR = 0 corresponds to very poor quality. Since increasing the minimum SNR would reduce the dynamic range that could be reported in marketing material, there is little incentive to change it.

Key points

- Camera dynamic range is *not* the same as image sensor dynamic range. For HDR image sensors, which have dynamic ranges specified at 120 dB or greater, camera dynamic range is often much lower (< 90 dB for the best cameras we've seen).
- Flare light degrades dynamic range by fogging shadows, but there could be cases where medium-range flare light (which falls off with distance from light sources) makes traditional dynamic range measurements look better than reality. This could potentially lead to a situation where increasing the flare improves the slope-based dynamic range.
- Because flare light affects *all* camera dynamic range measureements (traditional and new), it will be extremely important to standardize the chart design and capture environment.
- Slope-based dynamic range DR_{slope} yields large numbers that marketers like, but should not be used because it includes dark regions where SNR is unacceptably low, i.e., no image features are visible.
- The 95 dB (4.75 OD) tonal range of the Contrast Resolution chart is well below the dynamic range of the best HDR sensors. But we have *never* seen a camera that displayed detail in the bottom row of the chart, i.e., had a dynamic range over 80 dB. We had to photograph the chart with the upper rows masked out to prove to ourselves that detail was really present in the bottom row.

Future work

- Thresholds for human visibility need to be correlated to Contrast Resolution measurements using perceptual studies. Thresholds for machine vision detectability also need to be determined.
- The new dynamic range measurements need to be brought before appropriate standards organizations.
- Automotive companies use standards to protect themselves when accidents take place. Inadequate standards can cause a lot of trouble. It's best to get them right sooner rather than later. To that end we will continue working with the IEEE P2020 [6] Standard for Automotive System Image Quality group, particularly with the IQ Computer and Human Vision subgroups. We encourage readers to consider joining.

References

[1] A. Rizzi, J. McCann, "Glare-limited appearances in HDR images," Journal of the SID 17/1, 2009.

[2] E-V. Talvala, A. Adams, M. Horowitz, M. Levoy, "Veiling Glare in High Dynamic Range Imaging," ACM Transactions on Graphics, Vol. 26, No. 3, Article 37, July 2007.

[3] D. Wueller, "Image Flare measurement according to ISO 18844," IS&T Electronic Imaging conference, 2016.

[4] European Machine Vision Association, EMVA Standard 1288, [Online]. Available: www.emva.org/

[5] Samuel W. Hasinoff et al., "Burst photography for high dynamic range and low-light imaging on mobile cameras", ACM Trans. Graph., Vol. 35, No. 6, Article 192, November 2016.

[6] IEEE P2020 - Standard for Automotive System Image Quality [Online]. Available: https://standards.ieee.org/develop/project/2020.html

Author Biography

Norman Koren became interested in photography while growing up near the George Eastman House photographic museum in Rochester, NY. He received his BA in physics from Brown University (1965) and his Masters in physics from Wayne State University (1969). He worked in the computer storage industry simulating digital magnetic recording systems and channels from 1967-2001. He founded Imatest LLC in 2003 to develop software and test charts to measure the quality of digital imaging systems.