This is a preprint of
66926-41 paper in SPIE/IS&T Electronic Imaging Meeting, San Jose, January, 2007

Veiling glare: the dynamic range limit of HDR images

J. J. McCann*[a] & A. Rizzi[b],

[a]McCann Imaging, Belmont, MA 02478
[b] Dip. di Tecnologie dell'Informazione, Università degli Studi di Milano, Italy

Veiling glare: the dynamic range limit of HDR images

J. J. McCann*[a] & A. Rizzi[b],
[a]McCann Imaging, Belmont, MA 02478
[b] Dip. di Tecnologie dell'Informazione, Università degli Studi di Milano, Italy

ABSTRACT

High Dynamic Range (HDR) images are superior to conventional images. However, veiling glare is a physical limit to HDR image acquisition and display. We performed camera calibration experiments using a single test target with 40 luminance patches covering a luminance range of 18,619:1. Veiling glare is a scene-dependent physical limit of the camera and the lens. Multiple exposures cannot accurately reconstruct scene luminances beyond the veiling glare limit. Human observer experiments, using the same targets, showed that image-dependent intraocular scatter changes identical display luminances into different retinal luminances. Vision's contrast mechanism further distorts any correlation of scene luminance and appearance.

There must be reasons, other than accurate luminance, that explains the improvement in HDR images. The multiple exposure technique significantly improves digital quantization. The improved quantization allows displays to present better spatial information to humans. When human vision looks at high-dynamic range displays, it processes them using spatial comparisons.

Keywords: HDR imaging, veiling glare, calibration of multiple exposure techniques, spatial algorithms, Retinex, ACE

1. INTRODUCTION

Many labs are studying High Dynamic Range (HDR) image capture and display. In digital imaging, the popular multiple exposures technique[1] attempts to extend cameras' range in recording darker scene regions. This paper measures camera and human responses to calibrated HDR test targets. It describes measurements showing that image dynamic range is limited by the camera's veiling glare.[2] Glare is an uncontrolled spread of an image-dependent fraction of scene luminance caused by unwanted scattered light in the camera and in the eye. The goal of the paper is measure the physical limits of veiling glare and to evaluate the role of accurate reproduction of luminance in imaging.

Multiple exposure techniques for rendering (HDR) scenes go back to the earliest days of negative–positive photography.[3] H.P. Robinson's 1858 composite print "Fading Away" was made using five exposed negatives.[4] Mees's long career (1890-1956) studying photographic sensitometry at University College London and Kodak established standards for high-dynamic range image capture on the negative, and high-slope rendering on prints.[5] Mees's "The Fundamentals of Photography", 1920, shows an example of a print made with multiple negatives with different exposures.[6] Ansel Adams's zone system established three sets of procedures: first, for measuring scene radiances; second, for controlling negative exposure to capture the entire scene range, and third, spatial control of exposure to render the high-range negative into the low-range print.[7]

In 1968, Land demonstrated the first electronic (analog) HDR rendering in his Ives Medal Address to the OSA.[8] Here the intent was to render HDR images using spatial comparisons. In 1978 Frankle and McCann patented a very efficient digital spatial-comparison HDR algorithm. In 1984&5 McCann described HDR image capture using low-slope film[9] in Siggraph courses. In 1997 Debevec and Malic used multiple exposures and least-square fits to solve for camera response function and the luminance of each pixel in the scene.[1] In this paper the rendering intent is to accurately determine the scene luminance of each pixel for processing and display.

Comparisons of actual luminances versus camera-based estimates of luminance show that the estimates' accuracy is limited by veiling glare.[2] ISO 9358:1994 Standard, " Veiling glare of image forming systems"[10] defines veiling glare and describes two standard techniques for measuring it. It describes how veiling glare is the sum of individual stray light contributions to a single pixel from all other light from the scene, even from light beyond the field of view of the camera. Stray light reflected from lens surfaces, sensor surfaces and camera walls causes veiling glare. The ISO standard defines the glare spread function (GSF), which is a measure of the amount of stray light as a function of angle from a small very

bright light source. Veiling glare is measured by ISO9358:1994 as the fraction of a very large white surround scattered into a small opaque central spot. For commercial camera lenses veiling glare values are in the range of 1 to 10 %, depending on the lens and the aperture.

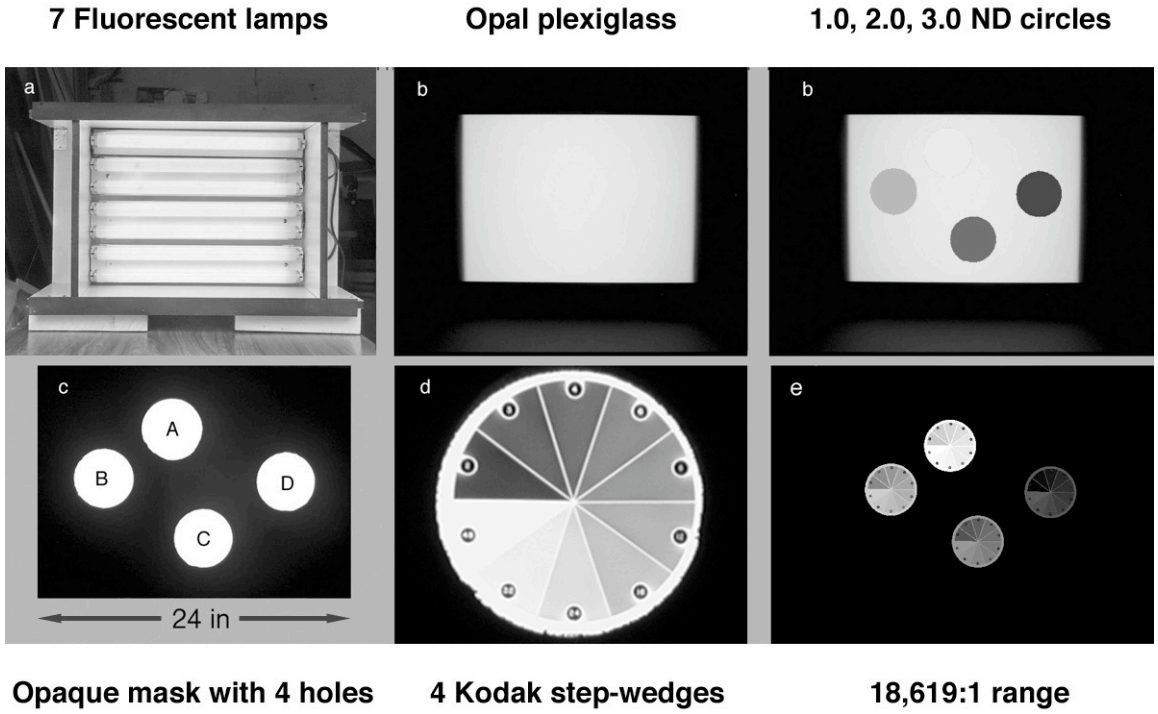

Figure 1a shows the light source made of 7 fluorescent tubes (20W). Figure 1b shows an opal-Plexiglas diffuser placed 6 inches in front of the lamps. Figure 1c shows the addition of 3 circular neutral density filters attached to the Plexiglas with densities of 1.0, 2.0, and 3.0. Figure 1d shows an opaque mask that covered the entire lightbox except for four circular holes registered with the N.D. filters. Figure 1e shows an enlarged view of a single Kodak Projection Print Scale. Figure 1f shows the assembled *4scaleBlack* target with a dynamic range of 18,619:1 [2049 to 0.11 cd/m^2]. Using opaque black masks, the luminance of each sector was measured with a spot luminance meter.

2. A CALIBRATED HDR DISPLAY

While ISO 9358:1994 provides a standard to compare different lenses and apertures, we wanted to measure the effects of veiling glare on HDR imaging. We used a single calibrated test target with 40 test luminance sectors (dynamic range = 18,619:1). Nearly 80% of the total target area was an adjustable surround; 20% of the area was luminance test patches. Using opaque masks to cover the surrounding portions of the scene, we photographed three sets of HDR test images with different amounts of glare. The experiment compares camera digits with measured scene luminance over a very wide range of luminances and exposure times. This experiment measured the extent that veiling glare distorts camera response in situations common with HDR practice.

In 1939, Kodak patented the Projection Screen Print Scale for making test prints[11]. It is a circular step wedge with 10 pie-shaped wedges, each with a different transmission. The range of transmissions was 20:1. After focusing a negative in an enlarger on the print film plane, dark-room technicians would place this scale on top of the unexposed print film in the dark. The wedges transmitted 82%, 61%, 46%, 33%, 25%, 17%, 14%, 9%, 8% and 4% of incident light so as to make a quick and accurate test print to select the optimal print exposure. Each pie-shaped segment gives a different exposure to each segment. The idea was to make a test-print under this scale so as to select the ideal exposure.

The components of our test display are shown in Fig 1. The display is made of transparent films attached to a high-luminance light-box. The are four Kodak Projection Print Scale transparencies mounted on top of 0.0 (ScaleA), 1.0 (B), 2.0 (C), and 3.0 (D) N.D. filters. The 40 test sectors are constant for both minimal (*4scaleBlack*) and maximal

(*4scaleWhite*) glare so that both targets have the same range of 18,619:1. For minimal glare, we covered all parts of the display except for the pie-shaped projection scales with an opaque black mask (*4scaleBlack*). For maximal glare, the opaque black mask was removed so that the zero-glare surround was replaced with maximal glare (*4scaleWhite*). The diagonal line in *4scaleWhite* is an opaque strip in front of the display.

3. CAMERA RESPONSE TO HDR TEST TARGET

We used the HDR calibrated target to measure the camera response with all three targets. With the *1scaleBlack* target we measure the camera response using only the 20:1 scale with lowest luminance (*1scaleBlack*). With the *4scaleBlack* target we measure the camera response using a display range of 18,619:1 with minimal glare. With the *4scaleWhite* target we measure the camera response using the same display range with maximal glare. We made separate sets of measurements with a digital camera and with a 35mm film camera using slope 1.0 duplication and with conventional negative films.

3.1 Digital Camera Response

We made photographs using a typical compact, high-quality digital camera (Nikon Coolpix 990) with manual, mid-range aperture (f 7.3) and exposure time controls. The experiment photographed 3 sets of images shown in Figure 2.

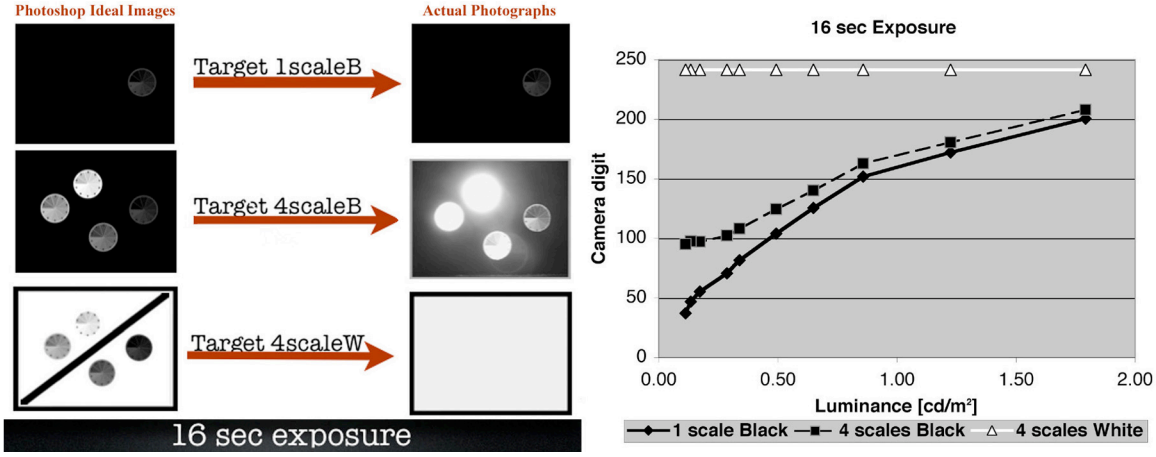

Figure 2 compares camera responses to the three displays using a single exposure time. This 16 sec time is the optimal exposure for *1scale B*. Figure 2(left) shows ideal synthetic Photoshop images of the three test targets and actual 16 sec exposures of the test target. The left column shows synthetic images in which digits, proportional to spot meter luminances, were pasted into each pie-shaped sector. The digital value for each sector was calculated using [256*Actual luminance/18,619]. These are the goal images with accurate 8-bit renditions of these scenes. The top row shows the goal image and an actual 16 sec exposure (best exposure for the lowest luminance scale D). The middle row shows the goal image and an actual 16 sec exposure of *4scaleBlack*. The bottom row shows the goal image and an actual 16 sec exposure of *4scaleWhite*.

Figure 2 (right) plots the surround's effect on camera digits. It plots the digits from ScaleD with 16 sec exposures. Diamonds plot the lowest glare data (*1scaleBlack*) for the 10 luminances. Squares plot *4scaleBlack* digits for the same luminances when the higher luminance scales [A, B, C] are unmasked. Glare from the other test sectors has increased camera digits. The darkest sector digit increased from digit 40 to 98. The triangles plot *4scaleWhite* digits with white surround. All camera responses equal the highest scanner digit of 242. The white surround is 77% of the total area of the image. Glare from the uncovered white surround had increased flux on the sensor to saturate the entire image. The three curves measure different camera responses to identical luminances.

Figure 2 (left) shows ideal synthetic Photoshop images of the three test targets and actual 16 sec exposures of the test target. The 16 sec exposure is optimal for recording the luminances of the lowest luminance scale D. Figure 2 (right) plots the camera digits for 10 Scale D luminances (0.11 to 1.79 cd/m^2) from all three photographs. The punctual luminance values at each wedge sector remain unchanged in the scene. A 16 sec exposure of the *1scale Black* target shows a typical camera response with digits from 37 to 201. Veiling glare has a small, but significant incremental effect on camera response to *4 scales Black*. Veiling glare overwhelms the camera response to *4 scales White*. The intent of multiple exposures in HDR imaging is to synthesize a new image with a significantly greater dynamic range record of the scene. The idea is simply to assume that scene flux [(cd/m^2) * sec] generates a unique camera digit. Longer

exposures collect more scene photons, and can move a dark portion of the scene up onto a higher region of the camera response function. This higher exposure means that the darker portions of the scene have better digital segmentation, more digits for better quantization of luminances. The HDR multiple exposure proposal[1] claimed to extend the cameras range by calibrating flux (luminance * time) vs. camera digit. This assumption is correct up until unwanted veiling glare distorts the camera responses. The results in Figure 2 show that glare can be substantial.

We measured the influence of veiling glare using 16 scene exposures. Using a fixed aperture we photographed the same low-luminance scale with the three surrounds presented above using a series of different exposure times. The plots of camera digits resulting from calibrated scene luminances are shown in Figure 3.

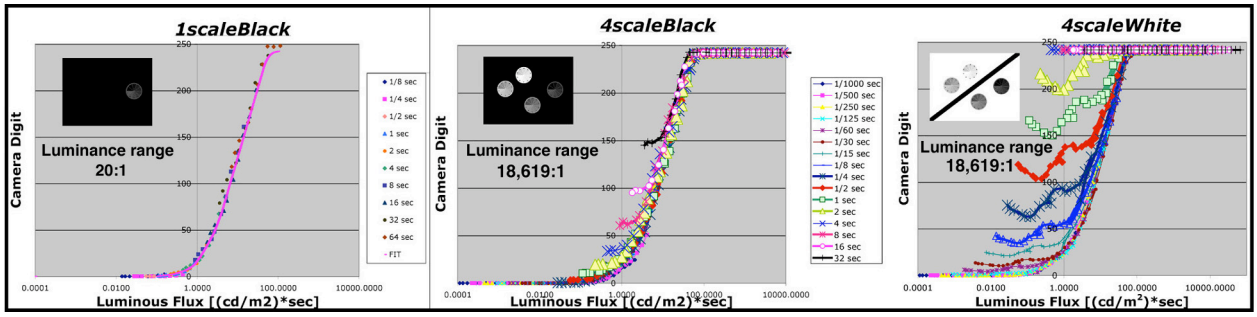

Figure 3 plots camera digits vs. luminous flux for all scales, for all exposures, for the three sets of photographs.

Figure 3(left) plots *1scaleBlack* (range 20:1) camera digits. It shows the desired coincidence of camera digits and flux. The curve provides us with an accurate camera response function. The sensor digits saturate at 247 with a flux of 78.4 sec*cd/m^2; at digit 1 the flux is 0.107 sec*cd/m^2. The camera dynamic range is 731:1, or 2.86 $\log_{10}$ units.

Figure 3(center) plots *4scaleBlack* (range 18,619:1) camera digits. It shows the minimal effects of glare for this range and configuration using a black surround between test scales. For an optimum exposure (1/2 sec) the sensor digits saturate at 242 with a flux of 119 sec*cd/m^2; at digit 11 (departure from sensor curve) the flux is 0.84 sec*cd/m^2. The glare limited dynamic range is 141:1 or 2.15 $\log_{10}$ units. The effects of glare are seen in low-luminance sectors.

Figure 3(right) plots *4scaleWhite* (range 18,619:1) camera digits. It shows the maximal effects of glare for this range and configuration with a white surround between test scales. Image dependent glare distorts luminance estimates for luminances. For an optimum exposure (1/15 sec) the sensor digits saturate at 241 with a flux of 71 sec*cd/m^2; at digit 40 (departure from sensor curve) the flux is 1.02 sec*cd/m^2. The glare limited dynamic range is 77:1 or 1.89 $\log_{10}$ units.

The *1scaleBlack* photographs have the lowest veiling glare and are an accurate measure of the camera sensor response function. The only sources of glare are the test patches themselves (range 20:1).

In *4scaleBlack* the camera's digit responses to four 10-step scales attempt to capture a combined dynamic range of 18,619:1. This target measures the minimum glare for a scene with this range, because it has an opaque black surround. The only source of glare is the test patches that vary from 2094 to 0.11 cd/m^2. The data shows that camera digit does not predict camera flux because the data fails to fall on a single function. The same digit is reported from different luminances. This is important because this display was intended to measure the minimal glare for an 18,619:1 image.

When we removed the black mask covering the lightbox in background we go to the situation with maximal veiling glare (*4scaleWhite*). Nearly 80% of the pixels are making highest possible contribution to veiling glare. Here the influence of glare is dramatic. Camera digits are controlled as much by glare as by luminance.

The data from all three sets of photographs are different. Data from *1scaleBlack* provides an accurate measure of camera sensor response vs. luminance – a single response function. Data from *4scaleBlack* shows variable camera sensor response vs. luminance. Data from *4scaleWhite* shows many scene-dependent responses vs. luminance. Camera digits from multiple exposures cannot provide a trustable means of measuring HDR scene flux. Camera digits cannot accurately record HDR scene flux because of glare. Veiling glare is image dependent. We also have performed tests using different cameras, and various changeable lenses, and we obtained similar results.

We used data from all Scale D *1scaleBlack* exposures to measure the camera response to flux. We compared the multiple-exposure technique flux estimates to actual flux for *4scaleBlack* and *4scaleWhite* targets. *4scaleBlack* has no glare from 77% of the image area, yet shows worst-case errors as large as 300% distortions. *4scaleWhite* (77% area with

maximal glare) shows 10,000% errors. All possible backgrounds with Scales A, B, C, D will fall between these data sets. There are different, large glare distortions for the same luminance depending on exposure. [2]

3.2 Film Camera Response

We made another set of photographs with a typical high-quality 35 mm film camera (NikonFM with a Nikkor 50mm 1:2 lens) using Kodak Slide duplication film. This follows the single exposure HDR capture technique described by McCann in tutorials at Siggraph conferences in 1984 and 1985. Slide duplication film has slope 1.0 on a log exposure vs. log luminance plot. In other words, output luminance equals input luminance. Since it is a color film it can be scanned for color and does not require calibration to remove the color masks found in color negative film. Here we use multiple exposures to capture both 18,619:1 displays (*4scaleBlack* and *4scaleWhite*). The 2 sets of scanned-duplication-film digits are shown in Figure 4.

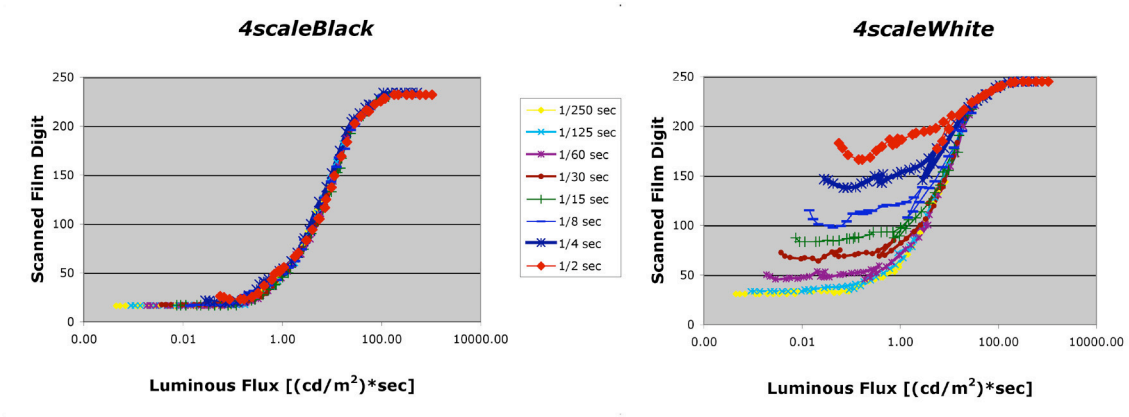

Figure 4 plots scanned film digit vs. luminous flux for 8 exposures ranging from 1/250 sec to 1/2 sec. The left graph plots data for *4scaleBlack*. This camera film and scanner system has less veiling glare than the digital camera in Figure 3. The data from the 8 different exposures superimpose to form a single function, except for the very bottom. The right graph plots data for *4scaleWhite*. Here the white surround adds veiling glare to generate 8 different response functions.

3.3 Negative Film Camera Response

We made another set of photographs with the same typical high-quality 35 mm film camera (NikonFM with a Nikkor 50mm 1:2 lens) using Kodak Max 200 negative film. Here we use single exposures to capture both 18,619:1 displays. We used 7 different exposures to measure the camera-film-scanner process using the low glare 20:1 single scale. The 3 sets of scanned negative film digits are shown in Figure 5.

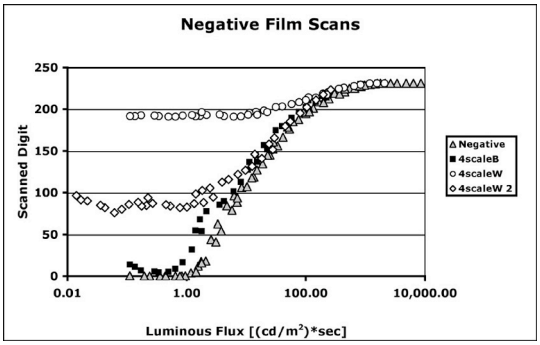

Figure 5 plots scanned single negative digits vs. $\log_{10}$ luminous flux for *4scaleBlack* (black squares) and *4scaleWhite* (white circles and diamonds) targets. The *1scaleBlack* (gray triangles) data plots data from 7 negatives with different exposures. The triangles report the response of the camera-film-scanner process with the lowest level of glare. It shows that the negative process can accurately record fluxes from 2639 to 0.24 sec*cd/m^2. The curve from *4scaleBlack* saturates at 1181 sec*cd/m^2 and inverts at 0.34 sec*cd/m^2. The inversion is caused by glare that limits usable range. There are two different exposures for *4scaleWhite* target; one is 8 times longer than the other. The data from *4scaleWhite* (white circles) saturates at 1,181 cd/m^2 and inverts at 6.17 cd/m^2. The curve from *4scaleWhite2* (white diamonds) has a maximum digit at 2094 cd/m^2 and inverts at 5.8 cd/m^2.

Figure 5 shows that the negative-camera-scanner process can accurately record fluxes from 2639 to 0.24 sec*cd/m^2 (dynamic range of 11,100:1, or 4.05 $\log_{10}$ units) [See Table 1]. The data from *4scaleBlack* show a small effect of glare from the 40 test areas. This glare reduces the dynamic range of the image to 3.5 $\log_{10}$ units. The glare from the white surround in *4scaleWhite* and *4scaleWhite2* reduces the dynamic range of the image measurements of 2.3 and 2.6 $\log_{10}$ units.

Target	Max Flux	Min Flux	Range	Log Range
Negative	**2094.0**	**0.19**	**11,021**	**4.04**
4scaleB	**2094.0**	**0.33**	**6,345**	**3.80**
4scaleW	**2094.0**	**6.10**	**343**	**2.54**
4scaleW 2	**261.8**	**0.77**	**340**	**2.53**

Table 1 shows the range limits of the negative film plotted in Figure 5. It lists the maximum flux below system saturation (Max Flux), the minimum flux above digit reversal from glare (Min Flux), the ratio (Range), and $\log_{10}$ ratio (Log Range).

Table 1 shows that the negative-camera-scanner process in the lowest glare condition is capable of capturing slightly greater than 4.0 $\log_{10}$ units. The data also show that system response to the 18,619:1 *4scaleBlack* test target has a veiling glare limit of 3.6 $\log_{10}$ units. Two different *4scaleWhite* exposures have very different curves, but almost the same glare limited Log Range measurements (2.3 and 2.6 $\log_{10}$ units).

The fact that the dynamic ranges for the two exposures of the *4scaleWhite* target are almost the same is important. Their response curves in figure 5 are very different. The *4scaleWhite* scanned digits have a max of 231 and a min of 191. The *4scaleWhite2* scanned digits have a max of 223 and a min of 94. The range of digits representing the scene is only of secondary importance. The range of digits describes the number of quantized levels used to represent the image. It controls discrimination, but does not control the dynamic range of the image. Too often the number of bits of quantization is confused with scene and image dynamic ranges. The number of bits can only describe quantization. As seen in the above results, both the scene and the camera image dynamic range must be calibrated independently.

Conventional negative film can capture a greater range of luminances than falls on the camera image plane from these targets. The dynamic range of a single exposure negative-film-scanner process exceeds the glare limited *4scaleBlack* image by 0.5 $\log_{10}$ units and glare limited *4scaleWhite* image by 1.6 $\log_{10}$ units. Multiple exposures with negative films serve no purpose. The glare-limited ranges of the camera and these HDR scenes are smaller that the film system range.

The above data suggests that in high, and in average glare, scenes the image dynamic range on the film plane is less than 3.0 $\log_{10}$ units. Only in special cases, very low-glare scenes, does the limit exceed 3.0. The data here show how well the designers of negative films did in optimizing the process. They selected the size distribution of silver halide grains to make the negative have a specific dynamic range, around 4.0 $\log_{10}$ units. Thus, single-exposure negatives capture the entire range possible in cameras, with low glare scenes. For most scenes this image capture range provides a substantial exposure latitude, or margin of exposure error. After reading the papers of C. K. Mees, L. A. Jones, and H. R. Condit, it is easy to believe that this fact is not a coincidence. [12, 13,14]

4. HUMAN VISION RESPONSE TO HDR DISPLAY

The second effect of veiling glare on HDR imaging is intraocular scatter that controls the dynamic range of luminances on the retina. In section 3.0 we saw that camera lenses limit the range of luminances falling on the camera sensor plane. Human intraocular scatter limits dynamic range more than glass lenses. Here we will describe the range of discrimination and the corresponding range of retinal luminances. In addition, we will measure observed appearance for both *4scales Black* and *4scales White* test targets.

4.1 Veiling Glare on the Retina

In 1983 Stiehl et. al.[15] described an algorithm that calculated the luminance of each pixel in an image on the human retina after intraocular scatter. It calculated the luminance at each pixel on the retina based on that display pixel's measured luminance and the calculated scattered light from all other scene pixels. The calculation used Vos and Walraven's measurements of the human point spread function. Stiehl measured the actual display luminance and calculated the retinal luminance for the display shown in Figure 6.

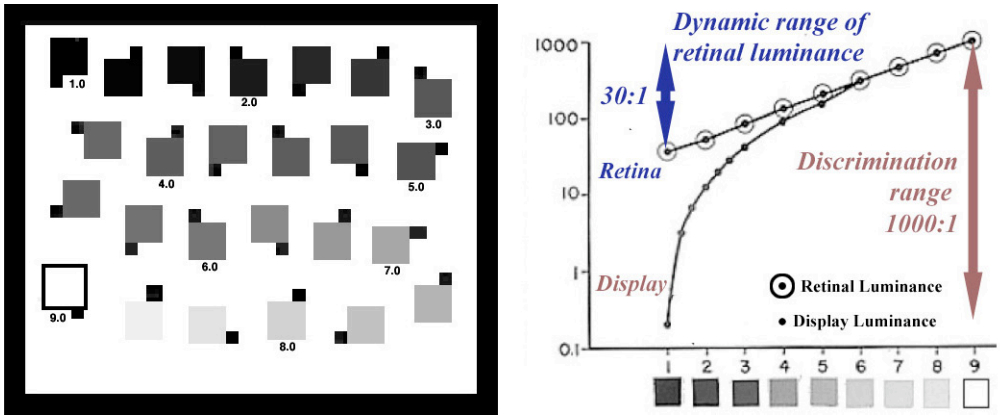

Figure 6 (left) shows an experimental lightness scale chosen by observers to have equal steps in appearance (bisection). Observers selected transparent filters so that the 9 steps were equal steps in apparent lightness. In Figure 5 (right) dots plot the display luminance for the HDR transparent lightness scale. The circled dots plot the luminance on the retina, after intraocular scatter. Observers can discriminate display luminances over of 3 log units. This discrimination is made using calculated retinal luminances of only 1.5 log units.

These results show that measures of discrimination are completely distinct from measurements of dynamic range. Humans continue to discriminate appearances of display blacks that are 1/1000th the white luminance, although the stimulus on the retina is limited by scatter to only 1/30th the white. Discrimination has to do with spatial comparisons. The fact that humans can discriminate luminance differences cannot be used in any analysis of dynamic range. Discrimination depends on the local stimulus on the retina.[16] Discrimination and dynamic range are scientifically unrelated.

4.2 Visual Appearance of HDR Displays

We asked observers to evaluate the appearance of the *4scaleBlack* and *4scaleWhite* displays using magnitude estimation. Observers sat 48 inches from the 24 inches wide display. The radius of each sector was 2 inches; subtending 2.4 degrees. Three observers were asked to assign 100 to the "whitest" area in the field of view, and 0 to the "blackest". We then instructed them to find a sector that appeared middle gray and assign it the estimate 50. We then asked them to find sectors having 25 and 75 estimates. Using this as a framework the observers assigned estimates to all 40 sectors. The data from each observer (ages 31, 64, 68) was analyzed separately. No difference between observers was found. The average results are shown in Figure 7.

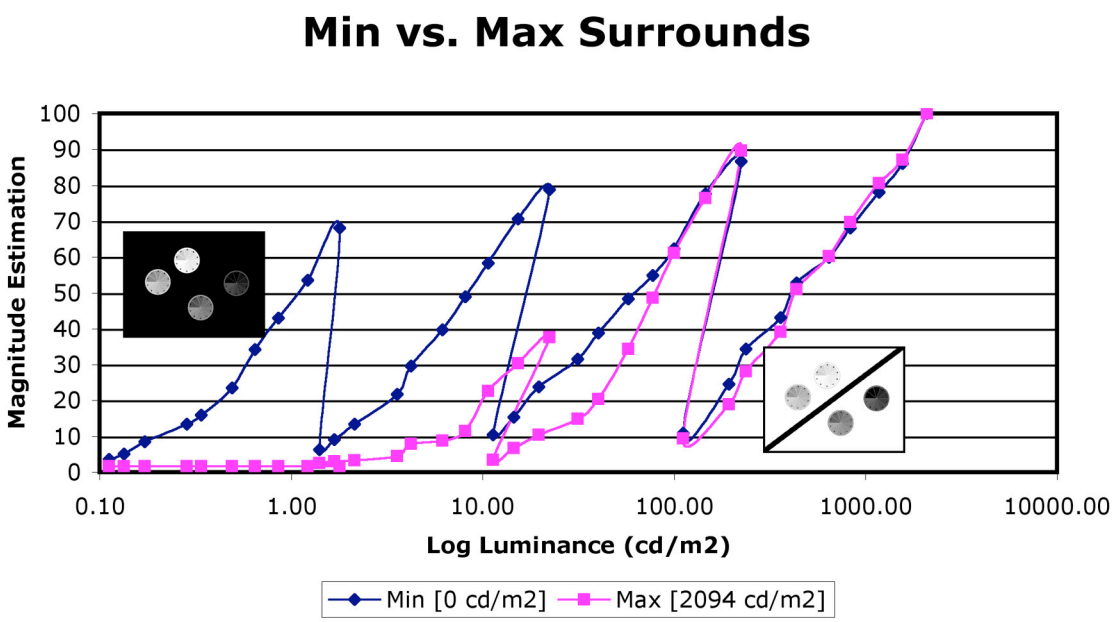

Figure 7 plots magnitude estimation of appearance vs. calibrated luminance for the 40 sectors in *4scales Black* and *4scales White* test targets. Although the luminances are exactly equal the appearances are not. With a black surround observers can discriminate all 10 sectors in all four displays. With a white surround observers cannot discriminate below 2 cd/m^2.

The results of Figure 7 show very dramatically the role of spatial comparison and scattered light in vision. The *4scales Black* and *4scalesWhite* appearance estimates overlap for the top 5 luminances. Below that, contrast makes the luminances in the white surround darker. The white surround makes the local maxima in scales C and D darker than in the zero-luminance surround. Scattered light from the white surround severely limits the discrimination below 10 cd/m^2. The *4scalesBlack* estimates are very different. Here, observers can discriminate all 40 test sectors. As shown in an earlier paper in the symposium[17] the maxima generate appearances that fall between constancy and luminance dependence. They fall on a low-slope Hipparchus line analogous to stellar magnitude. The data in figure 7 shows that these results hold for local, as well as for global, maxima. The estimates for sectors less than the local maxima decrease with a higher slope function.

Image dependent intraocular scatter transforms the identical display luminances into different sets of retinal luminances. Contrast mechanisms using spatial comparisons further distort any correlation of scene luminance and appearance. Any attempt to find an appropriate "tone scale" relating camera digit to appearance is futile. HDR does not accurately render luminances, and appearances are controlled by spatial comparisons with other pixels in the image. Any function that uses a single pixel's luminance as input cannot account for both the physical image-dependent effects of scatter and the physiological effects of spatial comparisons in contrast.

5. DISCUSSION

Veiling glare limits HDR imaging in two distinct ways. First, camera glare limits the range of luminances that can be accurately measured (Section 3). Multiple exposures improve the quantization of digital records, but fail to accurately record scene data. Second, intraocular scatter limits the range of scene luminances falling on the retina (Section 4). Scatter is the physical basis of cube-root appearance scales such as L* in L*a*b*.[15]

Accurate camera estimates of scene luminance are impossible for the 4.3 $\log_{10}$ dynamic range images studied here. The comparison of white and black surrounds shows a dramatic scene dependence with a single camera, using a constant aperture. In addition, the camera flux estimates, when compared with actual flux, show a different error with each exposure. It may be tempting to look for some type of average-flux curve that represents data with smaller errors, but that idea is in conflict with the fundamental aim of the process, namely recording accurate scene luminance. Multiple exposure HDR is limited by veiling glare that is scene-, exposure-, lens-, aperture-, and camera-dependent. The accuracy of scene-luminances estimates varies with all these parameters.

Some HDR algorithms attempt to correct for glare.[18] Given the characteristics of the camera, calculate the luminances in the scene. The glare spread functions of commercial lenses fall off very rapidly with distance to a very small value. We might think that such small glare values cannot affect distant pixels. However, there are millions of pixels that contribute glare to all other pixels. Each pixel is the sum of scene luminance plus scattered light from all other pixels. The sum of a very large number of small contributions is a large number. Sorting out these millions of scene-dependent contributions would be required to accurately correct for glare. ISO 9358:1994 Standard states unequivocally that: "the reverse [deriving luminance from camera response] calculation is not possible"[10].

Claims are made that recent multiple-exposure HDR algorithms capture wider scene luminances, or colors than previously possible.[19] These claims are severely limited by scene and camera veiling glare. Both camera and intraocular glare are image dependent and cannot be rigorously removed by calculation. As shown above, the designers of negative films selected a 4.0 $\log_{10}$ response range. That range exceeds the camera glare limit for almost all scenes. Further, the range of conventional transparency film (3.0 $\log_{10}$) equals the range of discriminable luminances in a display with a white surround. Nevertheless, the resulting HDR images are considerably better than conventional images with limited range. Since HDR imaging works so well, there must be reasons, other than accurate luminance, that explains the improved images. The multiple exposure technique does significantly improve digital quantization.

Veiling glare for human vision is much worse than for cameras.[15] Nevertheless, human vision has a much greater apparent dynamic range than camera systems. Humans can see details in highlights and shadows much better than conventional films and electronic cameras can record. Camera response functions are tuned for low-contrast, uniform-illumination scenes. They generate high-contrast renditions of low-range scenes. Early HDR algorithms[20] never attempted to determine actual scene luminance, since luminance is almost irrelevant to appearance.[21] Instead, these spatial algorithms mimicked vision by synthesizing HDR visual renditions of scenes using spatial comparisons. The intent of Land and McCann's electronic HDR imaging was to render high-dynamic range scenes in a manner similar to human vision. They made the case for spatial comparisons as the basis of HDR rendering in the B&W Mondrian experiment.[22] There, a white paper in dim illumination had the same luminance as a black paper in high illumination, but their appearances were strongly different.

Gatta's thesis[23] reviews many papers that combine HDR capture with a variety of tone-scale rendition functions. Tone scales cannot improve the rendition of the black and the white areas in the Mondrian with the same luminance. Tone-scale adjustments designed to improve the rendering of the black, do so at the expense of the *white in the shade*. As well, improvements to white make the *black in the sun* worse. When two Mondrian areas have the same luminance, tone-scale manipulation cannot improve the rendering of both white and black. Land and McCann made the case that spatial algorithms can automatically perform spatial rendering shown by Adams to compress HDR scenes into the limited range of prints. Such rendering is not possible with tone-scale manipulations in which all pixels are modified by a single tone–scale function.

The true benefit of high-dynamic-range image capture is improved quantization that can be used in spatial comparison algorithms. Spatial comparison images correlate with appearance, but not with scene luminance. By preserving the original scene's edge information, observers can see details in the shadows that are lost in conventional imaging. Spatial techniques have been used by painters since the Renaissance, multiple exposures and dodging and burning have been used by photographers for 150 years, and digital spatial algorithms, such as Retinex[8] and ACE[25] have been used to display high-range scenes with low-range media. HDR imaging is successful because it preserves local spatial details. This approach has shown considerable success in experimental algorithms [24,25,26] and in commercial products.[27]

6.0 CONCLUSIONS

This paper measures how much veiling glare limits HDR imaging in image capture and display. Glare is the scene- and camera- dependent scattered light falling on image sensors. First, glare limits the range of luminances that can be accurately measured by a camera, despite multiple exposure techniques. We used 4.3 $\log_{10}$ dynamic range test targets and a variety of digital and film cameras. In each case, the camera response to constant luminances varied considerably with changes in the surrounding pixels. HDR image capture cannot accurately record the luminances in these targets. Second, we measured the appearance of the same targets. Appearance did not correlate with luminance: it depended on physical glare and physiological contrast. The improvement in HDR images, compared to conventional photography, does not correlate with accurate luminance capture and display. The improvement in HDR images is due to better digital quantization and the preservation of relative spatial information.

7.0 ACKNOWLEDGEMENTS

The authors wish to thank S. Fantone, D. Orband and M. McCann for their very helpful discussions.

REFERENCES

*mccanns@tiac.net, rizzi@dti.unimi.it

1 P. E. Debevec, J. Malic, "Recovering high-dynamic range radiance maps from photographs", *ACM SIGGRAPH*, 369, 1997.
2 J. J. McCann and A. Rizzi, *"* Spatial Comparisons: The Antidote to Veiling Glare Limitations in HDR Images", in *Proc ADEAC/SID&VESA 3rd Americas Display Eng. & App. Conference*, Atlanta, 155-158, 2006
3 J. J. McCann and A. Rizzi, "Spatial Comparisons: The Antidote to Veiling Glare Limitations in Image Capture and Display" in *IMQA2007 The Second Int. Workshop on Image Media Quality and its Applications*, March, 2007, Chiba, Japan, in press.
4 T.Mulligan & D.Wooters, eds, *Photography from 1839 to today, George Eastman House, Rochester, NY*, Taschen, Koln, 360, 1999.
5 C. E. K. Mees, *An Address to the Senior Staff of the Kodak Research Laboratories*, *Kodak Research Laboratory*, Rochester, 1956.
6 C. E. K. Mees, *The Fundamentals of Photography*, *Kodak Research Laboratory*, Rochester, p.82, 1920.
7 A. Adams, *The Negative*, New York Graphical Society, Little, Brown & Company, Boston, pp.47-97, 1981.
8 E. Land and J. J. McCann, "Lightness and Retinex Theory", *J. Opt. Soc. Am*., **61**, pp. 1-11, 1971.
9 J. J. McCann, "Calculated Color Sensations applied to Color Image Reproduction", in *Image Processing Analysis Measurement and Quality*, Proc. SPIE, Bellingham WA, **901**, pp. 205-214, 1988.
10 "ISO 9358:1994 Standard, "Optics and optical instruments. Veiling glare of image forming systems. Definitions and methods of measurement", (ISO, 1994).
11 J. W. Gillon, U.S. Patent, 2,226,167, December 4, 1940.
12 C. E. K. Mees, *Photography*, The MacMillan Company, New York 1937.
13 L. A. Jones and H. R. Condit, *J. Opt. Soc. Am*. **38,** 1-11, 1948.
14 C. E. K. Mees and T.H. James, *The Theory of the Photographic Process*, 3rd ed., The MacMillan Company, New York 1966.
15 W.A. Stiehl, J. J. McCann, R.L. Savoy, "Influence of intraocular scattered light on lightness-scaling experiments", *J. opt Soc .Am*., **73**, 1143-1148, 1983.
16 T. Cornsweet and D. Teller, "Relation of increment Thresholds to Brightness and Luminance", *J. Opt. Soc. Am*. **55,** 1303-1308**,** 1965.
17 J. J. McCann, "Aperture and Object Mode Appearances in Images", in Human Vision and Electronic Imaging XII, eds. B. Rogowitz, T. Pappas, S. Daly, Proc. SPIE, Bellingham WA, **6292-26**, 2007.
18 E. Reinhard, G. Ward, S. Pattanaik, P. Debevec, *High Dynamic Range Imaging Acquisition, display and Image-Based Lighting*, Elsevier, Morgan Kaufmann, Amsterdam, chap 4, 2006.
19 http://www.cs.nott.ac.uk/%7Eqiu/jvci/Call-for-Papers.html
20 J. J. McCann, "Capturing a Black Cat in Shade: The Past and Present of Retinex Color Appearance Models", *J. Electronic Img*., 13, 36-47, 2004.
21 G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods Quantitative Data and Formulae*, 2nd Ed, Wiley, New York, chap 6, (1982).
22 E. H. Land and J. J. McCann "Lightness and Retinex Theory", *J. Opt. Soc. Am*., **61,** 1-11, 1971.
23 C. Gatta "Human Visual System Color Perception Models and Applications to Computer Graphics", PhD thesis, Università degli Studi di Milano, Dottorato di Ricerca in Informatica, XVIII ciclo, 2006.
24 J.J. McCann, ed., *Retinex at Forty*, Journal of Electronic Imaging, **13**, 1-145, 2004.
25 A. Rizzi, C. Gatta, D. Marini, "A New Algorithm for Unsupervised Global and Local Color Correction", *Pattern Recognition Letters*, Vol. 24, No. 11, July 2003, pp. 1663-1677.
26 J. J McCann, "Rendering High-Dynamic Range Images: Algorithms that Mimic Human Vision", in *Proc. AMOS Technical Conference*, Maui, p 19-28, 2005.
27 HP cameras with Digital Flash or Adaptive Lighting (Models 945, 707,717,727) use spatial comparisons to improve image rendering.