

Glare-limited appearances in HDR images

Alessandro Rizzi
John J. McCann

Abstract — Intraocular glare and simultaneous contrast control appearance in high-dynamic-range (HDR) images. This paper describes unique test targets that simulate real images. These targets change the HDR range by 500 times, without significantly changing the veiling glare on the retina. As well, these targets have nearly constant simultaneous contrast. The range of appearances possible from HDR images with different average luminances were measured. The targets displayed a maximum luminance range of 5.4 log units. Using magnitude estimates (MagEst) of appearances, the relationship between luminance and lightness from white to black was measured. With one exception, only small changes in appearance with large changes in dynamic range were found. It was also found that appearance was scene-dependent. The same dark grays (MagEst = 10) were observed with luminances of 10, 4.2, 1.1, and 0.063, depending on the percentage of white area in the surround. Glare from more white increases the retinal luminance of the test areas. Simultaneous contrast counteracts glare by making the appearance range (white–black) with a much smaller range of luminances. Appearance is controlled by both the optical scattered light and the spatial processing. A single tone-scale function of luminance cannot describe appearance controlled by scatter and spatial processing.

Keywords — HDR imaging, range of appearance, magnitude estimation, contrast, glare.

DOI # 10.1889/JSID17.1.3

1 Introduction

High-dynamic-range (HDR) imaging has four distinct luminance ranges that each play a different role in image rendering and in image reproduction. They are the luminance range in the scene, the captured luminance range in the image recorder, the display luminance range, and the range of luminances usable by humans in seeing the displayed image. First, we need to describe the image-making intent. At times, the goal is to *render* the scene as it appears. At other times, the goal is to *record and reproduce* the scene luminances. The most important lesson about image-making is that the best renderings of appearance do not reproduce luminances.¹ In principle, correctly reproducing the luminance value at every point in the entire scene must give to the observer the same stimulus to achieve photorealism. If two fields of view are identical, then their appearances have to be identical. In practice, this is impossible for the major part of real scenes.² Technical failures to capture all the information destroy appearance-matching, particularly if values near white are compressed.

At each step from scene, to capture, to display, to appearance there are limits to the dynamic range of information that is transferred to the next step. Camera lenses introduce veiling glare that reduces the range on the image plane.² Signal-to-noise characteristic of sensors limit the range of luminance response. Digitizing an image limits the range of the number of discrete luminance values. Too often, the number of quantization levels is confused with sensor range. Digitally speaking, cameras, displays, and

computers have a real cost for the bit-depth for each pixel. We can use a very large number of digital bits for the captured and displayed information, but this will not affect the size of the dynamic range of the sensor and/or the display. It makes economic sense that we use only the digital bit-depth precision for captured and displayed information that can be used by humans.

Scenes in the world can vary from the radiance of the sun to no light. This is the highest possible dynamic-range scene. Although most images do not include the sun in their field of view, many scenes have dynamic ranges that exceed three log units. The common myth that photography reproduces the radiance of scenes leads to the mistaken assumption that the quality of a reproduction is determined by the accuracy of reproducing scene radiances. Since the renaissance, artists have rendered HDR scenes successfully in low-dynamic-range media.^{1,3} Obviously, the range of the media used in rendering real-world scenes is much less than the sun/no-light range. Also, the range of a display media should take into account the range of information visible to humans. While surface reflections limit any reflective media's dynamic range, transparencies viewed in a dark room can have any range. Furthermore, since the film transmission is determined by the amount of dye formed, the transparency can have extremely low transmissions. In practice, it is much more efficient to use only the amount of dye in the image as is useful to the observer. Studies, by Mees, Pledge, Jones, Condit, and many others,³ set standards for transparency films so as to display the widest range of usable visible information.

Revised extended version of a paper presented at the 15th Color Imaging Conference (CIC 15) held November 5–9, 2007, in Albuquerque, New Mexico, U.S.A.

A. Rizzi is with the Università degli Studi Milano, Dipartimento di Tecnologie dell'Informazione, Via Bramante, 65, 26013 Crema (CR), Milan 26013, Italy; telephone +39-0373-898059, fax –898059, e-mail: rizzi@dti.unimi.it.

J. J. McCann is with McCann Imaging, U.S.A.

© Copyright 2009 Society for Information Display 1071-0922/09/1701-0003\$1.00

Digital HDR imaging research has devoted considerable study to the use of tone-scale maps that render HDR luminance and color data.⁴ Traditionally, the tone scale curve is the plot of log scene luminance *vs.* log final image luminance, as defined by Mees in 1920.³ It is the plot of the light input *vs.* output for each pixel for the entire image. We will use the traditional definition. Assuming that multiple exposures capture a wide range of scene radiances (camera flux digits),⁵ then the selection of a tone-scale map is one of the possible ways to scale the wide range of flux into the display device range. After considering the glare limitations tested in this paper, we suggest an alternative approach.

HDR image formats⁴ have been documented that encode dynamic ranges as large as 10^{76} . Even though some encode remarkably high dynamic ranges, the majority encodes around 3–4 log units. Recent papers published in this journal^{1,2} have shown that the range of luminances measurable in cameras is scene dependent. For most scenes the maximum range is between 3 and 4 log units. In this paper, we manufacture calibrated 5.4 log unit displays so as to measure of much dynamic range is useable by our visual system.

The human visual system (HVS) is, in fact, an optical system, and, as all optical systems, is subject to veiling-glare limitations. Glare is an uncontrolled spread of an image-dependent fraction of scene luminance caused by scattered light in the eye bulb by Tyndall scattering⁶ by macromolecules. Recent experiments have pointed out that veiling glare is a physical limit to HDR image acquisition.^{1,2,7–10}

In this paper, we measure the usable limits of luminance range for HDR displays usable by humans. We want to measure how veiling glare affects our visual tone-scale functions (luminance to appearance) in looking at HDR images. We do not need to store, and present in HDR displays, more luminance range than can be observed by humans. By limiting digital storage to the useful dynamic range, we can utilize more precise image quantization. Since bits, even if used in large numbers, are a finite resource, using limited dynamic ranges can result in a better quantization of perceivable tones. In displays, the expansion of dynamic range comes at a cost of technology. Using only the useful, visible dynamic range allows us to implement the best possible quantization in relation to the available disk space, color depth, and display technology. Whereas scenes can have sun/no-light dynamic ranges, technology limits the range available in the display of scenes. Displays of HDR scenes need only the range visible to humans.

2 Glare limits in HDR

Recently, to overtake the limited dynamic range of conventional displays, multiple exposure techniques⁵ have been combined with LED/LC displays that attempt to accurately reproduce scene luminances.¹¹ Veiling glare is not only a physical limit to HDR image acquisition, but it also limits the useful range of luminances in a display. Human intra-




FIGURE 1 — In a classic simultaneous contrast configuration, two opposing visual mechanisms contribute to the final appearance of the gray patches. Despite the fact that the gray square on the left (white surround) has higher retinal luminance from scatter, it looks darker. The HVS spatial processing more than cancels the effect of scatter.

ocular veiling glare determines the scene-dependent range of retinal luminance.^{3,12,13}

Human-observer experiments show two independent and opposing visual mechanisms. Intraocular veiling glare reduces the luminance range on the retina while physiological simultaneous contrast⁴ increases the apparent differences.^{1,2} Figure 1 shows the classic simultaneous contrast configuration. If we consider a gray patch surrounded by white, it will have much higher glare, due to the white surround. If retinal luminance predicts appearance, then it follows that this patch should appear lighter than the other on the black surround. However, simultaneous contrast, or human spatial image processing, makes the gray in white look darker. Glare distorts the luminances of the scene in one direction, and spatial contrast works to counteract glare.

To test how the veiling-glare limit can impact the HDR pipeline, we recently ran some experiments.^{1,2} We performed camera-calibration and human-observer experiments using a single test target with 40 luminance patches covering a luminance range of 18,619:1 (4.3 log units). In these experiments (Fig. 2), we measured the appearance of four identical transparent targets with four levels of illumination in the same scene in a black surround.² Observers measured appearance by making magnitude estimates (MagEst) between white and black. They were asked to assign 100 to the whitest areas and 1 to the blackest areas in the scene.

Average observer estimates are plotted in Fig. 2. The horizontal axis plots luminance measured with a spot photometer (cd/m^2). The vertical axis plots appearance (magnitude estimate value). The top target A has the highest luminance. It generates MagEsts from 100 to 11. The left target B, viewed through a 1.0 ND filter, has uniformly 10

²The term contrast has different definitions in photography and vision. In photography, it refers to the rate of change in reproduction luminance *vs.* scene luminance. It is the slope of the tone-scale function. In vision, it is the name of the spatial mechanism that enhances differences in appearance. A gray patch in a white surround is darker because of the physiology in the visual system, referred to as simultaneous contrast.




FIGURE 2 — The test target is on the bottom right. Four identical pie-shaped transparencies with ten different transmissions in pie-shaped areas were mounted on a light box. The top A had no neutral density filter behind it; B on the left had 1.0 ND filters; C on the bottom had 2.0 ND filters; D on the right had 3.0 ND filters. The surround was opaque. In total the target had 40 test areas with a luminance dynamic range of 18,619:1. The graph plots the average of observers' magnitude estimates of the appearance of the 40 test areas vs. luminance.

times less luminance than A. It generates MagEsts from 87 to 10. The bottom target C, viewed through a 2.0 ND filter, has uniformly 100 times less luminance than A. It generates MagEsts from 79 to 6. The right target D, viewed through a 3.0 ND filter, has uniformly 1000 times less luminance than A. It generates MagEsts from 68 to 4.

If we look along the horizontal line at MagEst = 50, we see that four different luminances (1.06, 8.4, 64, and 414 cd/m^2) generate the same appearance. If we look at luminance 147 cd/m^2 , we see that it generated both MagEst = 17 (near black) in A and MagEst = 87 (near white) in B. Similar examples of near-white and near-black appearances are found at luminance 15 (B&C) and 1.8 cd/m^2 (C&D).

Magnitude estimates of appearance in complex images do not correlate with luminance. In this target, nearly 80% of the total area is an adjustable surround; 20% of the area is luminance from the test patches. Removing the opaque background covering increased the glare to the maximum possible for this target configuration. With the white glare source replacing the opaque black, the observers' ability to estimate the patch magnitude strongly decreased (see Fig. 3). The range of discriminable patches decreased to less than 3 log units. In a black surround observers could discriminate all 40 luminance test areas over a range from 2049 to 0.11 cd/m^2 . When they replaced the black surround with a white (maximum glare) surround, the observers were unable to discriminate appearances below 2 cd/m^2 .

Vision's simultaneous contrast mechanism further distorted any correlation of scene luminance and appearance. In the black surround, lower luminances appeared much lighter than in the white surround. They showed that both physical intraocular scatter and the HVS contrast processing influenced the appearances of darker test targets.

McCann and Rizzi also measured the dynamic range of a camera-negative-film-scanner system on the same target.² The film was capable of recording 4.0 log 10 units of luminance. Glare from the 18,619:1 target surrounded by




FIGURE 3 — Removing the opaque background mask increases the glare to the maximum possible for this target (bottom right). Now the ability of the observers to estimate the patch magnitude strongly decreases to a range of less than 3 log units.

black reduced the range on the camera film plane to 3.5 log 10 units in a single exposure. The glare from a white surround further reduced the range to 2.4 log 10 units. The dynamic range of a single negative exposure exceeds the black surround scene (minimal glare) by 0.5 log units and white surround scene (maximal glare) by 1.6 log units.

3 Design of appearance scale target

We have two goals here. First, we want to measure the effect of doubling the dynamic range of the display. We want to do this using a surround pattern that holds constant both scattered light and spatial interactions, called simultaneous contrast. The second goal here is to measure the effect of similar targets with different amounts of scattered light and different spatial interactions.

What are the issues in designing a surround for measuring appearance? There are many papers that study how surrounds affect appearance.^{14,15} The spatial arrangement of luminances falls on the human spatial image-processing mechanisms and generates appearances that depend on size, separation,¹⁶ proportionality,¹⁷ and spatial-frequency distribution.¹⁸ Although we use the traditional name "simultaneous contrast" as the collective name for many phenomena, we need to keep in mind that spatial comparisons of retinal luminances are the underlying mechanism of seeing. We need to design our test target surrounds with both the spatial content and the luminance content in mind.

We want to measure the usable dynamic range of luminance using targets with a fixed amount of glare and with minimal changes in simultaneous contrast. To start, we set aside all the complexities introduced by gradients in illumination. We will just study patches of light that are uniform.

We could begin with luminance patches that are surrounded by no light. Bodmann¹⁹ showed that magnitude estimates of brightness plotted against luminance, in a black

surround, fit a low-slope line over 5 log units, similar to astronomer's standard for stellar magnitude.²⁰ Since these data come from small circles out of context on a black surround, they are inappropriate for typical scenes because they fail to account for the physical properties of scatter from normal images, as well as the HVS properties of simultaneous contrast. Appearance *vs.* luminance functions derived from experiments using black surrounds are different from those derived from complex images.

We could evaluate luminance patches in a white surround. Blacks appear blackest surrounded by white. However, veiling glare is greatest in white surrounds. The range of light, after scatter, from white surround luminances does not represent typical scenes that are made up of many different luminance areas. Appearance functions derived from experiments using white surrounds are also not appropriate. Nevertheless, we will measure appearance in a white surround as a control.

We could evaluate luminance patches in an average gray surround. Experiments compared lightness matches using a, white, gray, black, and complex-Mondrian surrounds. They showed that appearances in Mondrians are the same as those in a white surround, not gray surrounds.²¹ Gray surrounds show a rate of appearance against luminance between the low-slope black and the high-slope white. Appearance functions derived from experiments using average gray surrounds are also not appropriate.

If we consider the global physical properties of glare, we would like to have a surround that is, on average, equal to the middle of the dynamic range. This can be achieved by making the surround 50% max and 50% min luminance. Experiments have shown that the spatial distribution of white in the surround affects the appearance.¹⁸ To approximate real images, we distributed the half-white-half-black areas in different sized squares. Furthermore, we have energy over a wide range of spatial frequencies and can avoid the problem that simultaneous contrast depends on the size of the adjacent white areas.¹⁸




FIGURE 4 — Magnified view of two of 20 gray pairs of luminance patches. The left half (square A) has the same layout as the right (square B), rotated 90° counterclockwise. The gray areas in A have slightly different luminances, top and bottom. The gray areas in B have different luminances, left and right. The square surrounding areas are identical except for rotation. For each size there are equal numbers of min and max blocks.




FIGURE 5 — Target with 20 gray pairs of luminance patches. All gray pairs are close in luminance, but some edge ratios are larger than others.

Figures 4 and 5 shows the layout of our min/max test target. The display subtended $15.5^\circ \times 19.1^\circ$. It was divided into 20 squares, 3.4° on a side. Two 0.8° gray patches are within each square along with various sizes of max and min blocks. The two gray-square length subtends an angle approximately the diameter of the fovea. The smallest block (surrounding the gray patches) subtends 1.6 minutes of arc and is clearly visible to observers. Additional blocks 2x, 4x, 8x, 16x, 32x, and 64x are used in the surround for each gray pair.

4 Single- and double-density targets

The observers made magnitude estimates of the appearance of patches in single-density (SD) and double-density (DD) transparencies. The DD target is the aligned superposition of two identical photographic 4×5 -in. photographic (SD) films. Two transparencies double the optical densities. The whites in each transparency have an optical density (OD) of 0.19; the blacks have an OD of 2.89. The DD images have a min of 0.39 and a max of 5.78 OD (see Table 1). Both transparency configurations are backlit by four diffused neon bulbs.

Veiling glare for the HVS is a property of the luminance of each image pixel and the glare spread function (GSF) of the human optical system. Surrounds made up of

TABLE 1 — List of the luminances and optical densities of the min and max areas in SD and DD 50% white displays. Furthermore, it lists the display ranges and average luminances.

Target	Max cd/m ²	Min cd/m ²	Range:1	%Average cd/m ²	O.D. Min	O.D. Max	O.D. Range
Single Contrast	684	1.36	501	50.01%	0.19	2.89	2.70
Double Contrast	441	0.0018	251,189	50.00%	0.38	5.78	5.40

half-max and half-min luminances have very interesting glare properties for both SD and DD test targets. The average luminance of the SD target is 50.01% of the maximum luminance, from a display with a range of 501:1. The average luminance of the DD target is 50.00% of its maximum luminance, from a display with a range of 251,189:1. The effect of glare on the luminances of the gray test areas will be very nearly the same, despite the fact that the dynamic range has changed from 500:1 to 250,000:1. In other words, the black (min luminances) in both SD and DD targets are so low that they make only trivial contributions to glare. The white (max luminances) in both targets are almost equal and generate virtually all the glare. The layouts of both targets are constant, keeping simultaneous contrast stable. The physical contributions of glare are very nearly constant. By comparing the magnitude estimates of appearance of these SD and DD targets, we can measure the effects of constant glare on very different dynamic-range displays.

If the HVS can make use of the DD image range (250,000:1), then we expect to see a greater range of appearances in this image. If the veiling-glare limit has been reached in the SD image, then adding 500 times more range will have little or no effect on appearance.

5 Magnitude estimation experiments

The experiments were performed in a dark room. The only source of light was the target. The light-box had an average luminance of 1056 cd/m² (chromaticities $x = 0.45$, $y = 0.43$). Five observers made magnitude estimates of the appearance of the test patches between white and black. The observers were university students and workers between 18 and 23 years of age, with 20/20 or corrected 20/20 acuity. The five observers were asked to assign 100 to the “whitest” area in the field of view and 1 to the “blackest” appearance. We asked observers to use the same magnitude estimation scale for all test targets. Since the blackest “apparent black” is in a white surround, and since the series of experiments included black and near-black surrounds, we provided at the start of each experiment a white/black appearance calibration patch. On the side of the transparency was a 1-cm maximum-density square in the middle of a 3-cm minimum density square. We asked observers to use these white/black areas to assign extreme magnitude estimate values [100,1] if they did not appear in the target. These calibration paths were covered for the remainder of the experiment.

We instructed observers to find a square that appeared middle gray and assign it the estimate 50 (or very near value). We then asked them to find gray squares having 25 and 75 (or very near values) estimates. Using this as a framework, the observers assigned estimates to all squares (A-T in Fig. 5). Each of five observers repeated the experiment five times, not consecutively. They gave estimates for each half of the gray areas. We repeated the experiment with the same observers with SD and DD displays.

5.1 Magnitude estimation vs. Munsell lightness

There are a number of different appearance scales measured by asking observers to perform different tasks. Here, we ask observers to select a number equal to the magnitude of appearance. Other experiments, such as Munsell Lightness, asked observers to bisect the appearance white and black to measure middle gray. Further bisections measured the Munsell Lightness Scale. Stiehl *et al.*²² measured a bisection lightness scale using a HDR transparent display. The results were the same as those of Munsell, and the commonly used L^* in LAB. In a recent paper, we compared the Stiehl *et al.* lightness scale to the data described in this paper below. We found excellent agreement between Stiehl *et al.* and our magnitude estimates in 100% surround data.²³ In other words, for these experiments, we see no experimental difference between our magnitude estimation results, Stiehl *et al.* data, L^* and Munsell Lightness.

5.2 Average surround luminance = 50% max luminance

The first experiment measured the target shown in Fig. 5. The average results are shown in Fig. 6. They are the optimal tone-scale function for these 50% white complex scenes.

The plots for SD and DD nearly superimpose. The curves show the same asymptote at white and black. At middle gray (MagEst = 50), the SD curve is about 0.24 OD higher luminance. For the SD image, the highest luminance gray (area I) has a relative OD of 0.19, and an appearance estimate of 92. For that target the lowest luminance gray (area K) has an OD of 2.1 and an appearance estimate of 3.0. For the double-density image, area K has an OD of 4.1 and an appearance estimate of 1.8. The average of all observers on both targets show the same asymptote to black at OD of 2.3.




FIGURE 6 — Appearances of SD and DD displays with 50% average luminance surround. Observers estimate the same range limit of 2.3 log units.

These SD and DD images have nearly the same veiling glare and simultaneous contrast pattern. The curves are very similar, but do not overlap over most of their range.

The curves show the same asymptote at white and black. At middle gray (MagEst = 50), the SD curve is about 0.26 OD higher luminance.

The range of appearances from white to black is seen over 2.3 OD units. The results are consistent with veiling glare in determining the visible ranges. The effect of increasing the stimulus range has only a small effect because the SD image is at or near the maximum usable range on the retina for this scene.

5.3 Average surround luminance = 8% max luminance

The second experiment studied another pair of SD and DD targets with a different surround. We reduced the area of the white to 8% of the background, leaving the black to cover 92%. The effect of reducing the white area was to decrease the amount of veiling glare. The results are shown in Fig. 7.

Although different than the 50% white results, these curves are also similar over most of their range. The curves show the same asymptote at white and black. Again, at middle gray (MagEst = 50), the SD curve is about 0.30 OD higher luminance. The range of appearances from white to black is seen over 2.9 OD units. The results are consistent with veiling glare determining the visible ranges. The effect of increasing the stimulus range has only a small effect because the SD image is at or near the maximum usable range on the retina for this scene.

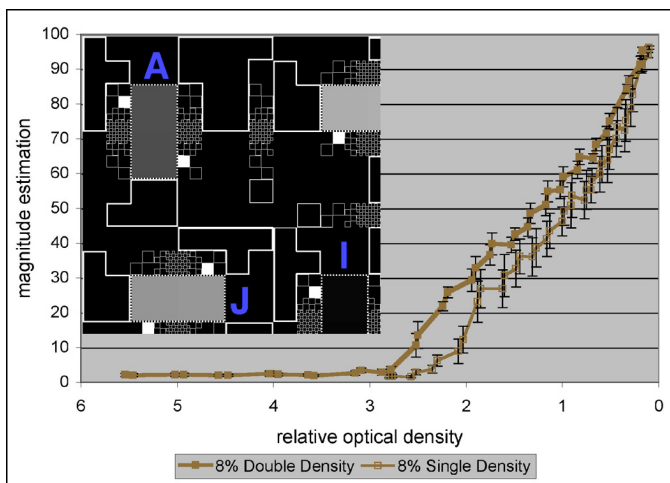


FIGURE 7 — Appearances of SD and DD displays with 8% average luminance surround. This figure shows the upper left corner of the 8% white target. Observers gave slightly different estimates for slope with the same range limit of 2.9 log units.




FIGURE 8 — Appearances of SD and DD displays with 100% average luminance surround. Observers gave slightly different estimates for slope with the same range limit of 2.0 log units.

5.4 Average surround luminance = 100% max luminance

The third experiment studied another pair of SD and DD targets with a 100% white surround. Increasing the white area increased the amount of veiling glare and the strength of simultaneous contrast. The results are shown in Fig. 8.

Again, these results differ from both the 50% and 8% white targets. The curves are similar over most of their range. The curves show the same asymptote at white and black. At middle gray (MagEst = 50), the SD curve is 0.23 OD higher luminance. The range of appearances from white to black is seen over 2.0 OD units. Again, the results are consistent with veiling glare determining the visible range.

5.5 Average surround luminance = 0% max luminance

The fourth experiment studied another pair of SD and DD targets with a 0% white surround. These displays have minimal veiling glare. The results are shown in Fig. 9.

For the first time, we see that the dynamic range of the display has a significant effect on appearance. The curves show the same asymptote at white, but diverge near black. At middle gray (MagEst = 50), the SD curve is 0.27 OD higher luminance. For very dark gray (MagEst = 5), the SD curve is about 1.28 OD higher luminance. The range of appearances from white to black is seen over 5.0 OD units.

Figure 9, with 0% white in the surround, shows minimal scatter. The retinal scattered light onto the gray patches is the smallest. Nevertheless, since all film transmittances are significantly lower in the DD image, the actual scatter increment added to gray-area luminance from the black surround is smaller in the DD image than in the SD image. Since the arrangement of gray squares is constant for both SD and DD displays, the scatter increment from grays is

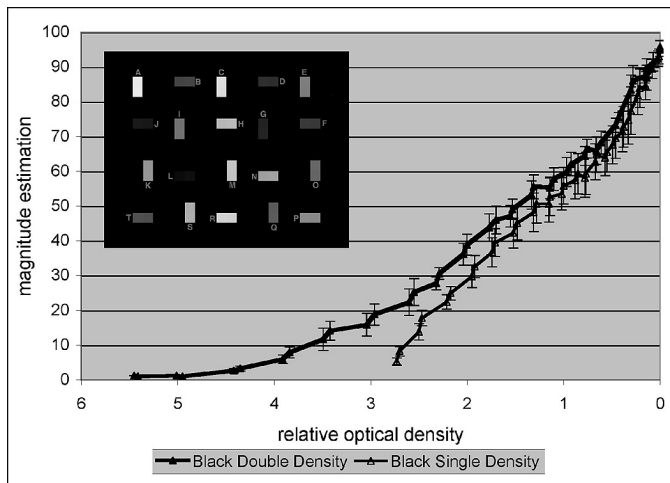


FIGURE 9 — Appearances of SD and DD displays with 0% average luminance surround. Observers gave slightly different estimates for slope with the same range limit of 5.0 log units.

also smaller in DD displays. With less scattered light, all the gray squares have less luminance in DD displays.²⁴

Now consider a gray square on a SD display that has the same target OD as a different square on a DD display. If appearance responded to retinal luminance, then MagEst values should be lower for the DD target because of the smaller scatter component of retinal luminance. In other words, the DD curve should fall to the right of the SD curve. As we see in Fig. 9, this luminance-based prediction does not correlate with appearance data. The DD curve falls to the left of the SD curve. We are led to conclude that the spatial-contrast mechanism not only cancels the effect of the decrease in intraocular scatter, but also increases the lightness appearance dramatically. For OD = 2.7, the MagEst is 5 in the SD target. The MagEst for the same target optical density in the DD target with less after-scatter retinal luminance is MagEst = 22. It has to be recalled that the same OD comes from different patches. In this case, the spatial-contrast mechanism has produced greatly different appearances from different minimal luminance surrounds. With 0% white surrounds, the effects of doubling the dynamic range can be observed. This effect does not correlate with retinal luminance. It correlates with spatial processing. Scatter would have shifted the curve to the right, and spatial processing cancelled that and moved the curve to the left and to higher magnitude estimates.

6 Discussion

Why should two different black surrounds have caused such very different response functions as seen in Fig. 9? The appearances of the black surrounds are the same. Why is it that two contrasts for the same spatial pattern are similar for MagEst near white and different MagEsts near black? Why is it that the DD target, with significantly lower luminances at every corresponding point, appears so much lighter than the SD target on the left half of Fig. 9?

The results seem surprising. One explanation of the surprise is the intellectual framework we choose to use in our thinking. Historically, simultaneous contrast has been used to illustrate an illusion, or a special case, in how humans see. It shows grays are about a 10% darker in white compared to black.²⁵ In this intellectual framework the role of the image's spatial content is secondary to the primary assumption that appearance is the result of luminance at each pixel. Land's Retinex Theory reversed these priorities. For Land, the appearance image is synthesized from the long-distance interactions of spatial comparisons. Appearance is neither the result of local processing nor global processing. It is the result of both. In this intellectual framework, there is only a small secondary role for luminance. Far more important are the values of the edge ratios, size, and distances between pixels, and from the maximum in the field of view.^{26,27} In 0% targets, there is only one small area of white. In SD and DD targets, corresponding gray areas are separated by the same distance from the white. However, the DD values of the edge ratios are double in OD. The expectation from spatial-contrast image processing is that such different targets, with very different luminances and edge ratios, should generate very different results. It is the spatial-contrast mechanism of the HVS that controls the appearance of dark grays on blacks in Fig. 9.

Looking at the data in Figs. 8, 7, and 6, we see somewhat analogous results. For the same OD in the two displays, the after-scatter retinal luminances for DD targets are lower than for SD targets. If appearance relies on retinal luminance, then DD targets should have a lower MagEsts for all squares. Using the luminance framework, we would expect DD data on the right side of the SD data. In all sets of data, we find the opposite.

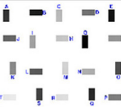
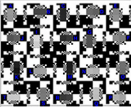
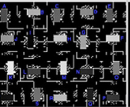
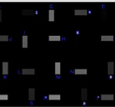
In all cases, the observers report very slightly higher appearances with double densities. We no longer see major changes in appearance. In the presence of white distributed throughout, the image we see small differences. These results are inconsistent with a luminance-based framework and consistent with a spatial-processing framework.

Figure 10 shows plots of four magnitude estimates of appearance for the DD targets as a function of target luminance. The slope of the transition from white to black depends on the amount of white in the background. Glare prevents the appreciation of most of the increase in dynamic-range information provided in the DD images.

The data in Fig. 10 shows that for four different backgrounds, there are four different optimal tone scales. Look at the horizontal line at MagEst = 10. Each of the four tone-scale plots intercepts the same dark-gray appearance at different luminances. The same dark gray (MagEst = 10) has luminances of 10, 4.2, 1.1, and 0.063, depending on the percentage of white area in the surround. With the same white in each test target, the luminance of the constant dark-gray appearance varies by a factor of 159:1.

The data shown in Table 2 shows the maximum usable range of luminance for each target design. Each one of the

TABLE 2 — Comparison of the percentage of white surround area with usable display dynamic range (log units).

% White	100%	50%	8%	0%
				
Usable Range (log units)	2.0	2.3	2.9	5.0

background configurations generates a different amount of glare (stable between SD and DD targets). Observer estimates show a usable range of 2.0 log units in the highest glare condition. Reducing the amount of white by one-half increases the usable range to 2.3 log units. We find a usable range of 2.9 with 8% white background. For the black background, observers can discriminate luminances over a 5 log unit range; this can be obtained only with a completely black surround and total darkness in the entire room. These very strict constraints are inconsistent with common scenes and normal viewing conditions.

The 100%, 50%, and 8% white displays held simultaneous contrast almost constant, while changing dynamic range. In the previous experiments described in Figs. 2 and 3 above, contrast and glare changed. Real scenes have variable amounts of simultaneous contrast and glare, and this combination presents a serious problem for tone-scale mapping. The Table 2 experiments used uniform illumination and constant local surrounds around each patch, so as to have nearly constant veiling glare. The results from Fig. 10 show that each image has its own tone-scale function, with a unique combination of glare and contrast.

The position of light sources and the placement of white areas in the scene control how glare plays a role on the

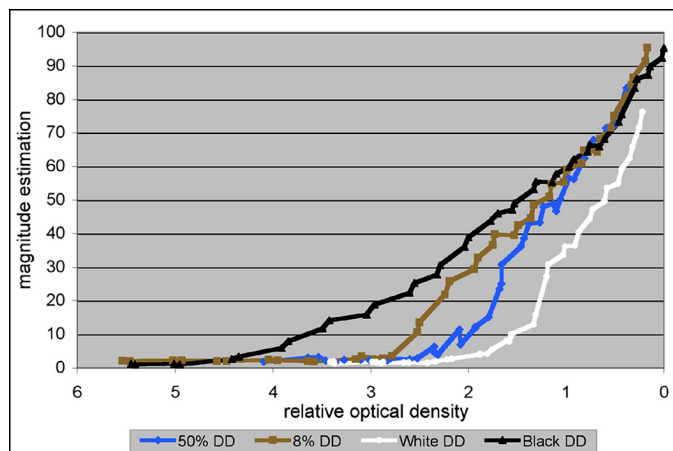


FIGURE 10 — Overall comparison of appearance slopes for DD displays with 100%, 50%, 8%, and 0%, average luminance surround. Observers measured significantly different slopes and dynamic-range limits.

image on the retina. In natural scenes, a big source of glare seldom has a huge vision field extension. As shown by Vos and van den Berg's²⁸ model of veiling glare, the amount of glare across the image changes quickly with distance. Nevertheless, there are finite amounts everywhere. Also, measurements of the effect of placement of white areas show that contrast mechanisms are nearly constant over large distances in the image.¹⁶ Both glare and contrast effect the observers' responses.

If we try to use the hypothesis that observer appearance can determine the best tone-scale map for luminance, we have a problem. For data in Fig. 10, this hypothesis requires a unique tone scale for each set of targets (two for 0% black). However, that unique tone scale varies according to the background and can only be calculated from spatial evaluations of the image incorporating corrections for both glare and HVS spatial-contrast processing. Calculating a scene-dependent tone-scale map would have to have two different components. First, it would have to calculate the scene luminance for the camera flux at each pixel. This would require identifying the percentage of glare for each pixel in the image. The ISO 9358:1994 standard states that the glare correction is impossible to compute from only the image luminances.²⁹ Second, the tone-scale map would have to calculate the image-dependent effects of HVS spatial processing for each pixel, as influenced by all other pixels. Recalling that Mees's definition of "tone scale" was explicitly an input-output mapping for a pixel, looking for ideal tone-scale maps for complex HDR images makes little sense.

Many natural HDR scenes have non-uniform illumination. How can HDR tone scaling predict the effects of non-uniform illumination? Land's Black and White Mondrian³⁰ studied illumination gradients. They presented a configuration where two areas had the same luminance and hence the same camera digit, *e.g.*, 128, but one was a white paper in dim light and the other a black paper in bright light. Since the two patches did not appear as equal, to improve the rendering in mapping the image dynamic we need to increase the digit for white and decrease it for black. This is impossible for a tone-scale curve to improve both whites and blacks since input 128 can have only one output value.

Humans are very good at discriminating very small increments in luminance at edges. As Cornsweet and Teller showed, the ability to discriminate depends on the local stimulus on the retina (after glare) and not on the appearance (where it is between white and black).³¹ Discrimination has to do with spatial comparisons. There is a long history of rendering HDR scenes that does not depend entirely on tone scales.^{1,9} Human vision, painting, and photography use spatial comparisons to synthesize a new low-range image from HDR input. Although the retinal receptors have a measured dynamic range of more than 10^{10} , the retinal ganglion cells transmitting information to the visual cortex have a range only slightly greater than 10^2 . Surface reflections from paintings and photographic prints limit their

range to less than 10^2 . Early electronic HDR algorithms synthesize new low-range images from HDR input.³⁰ The unifying principle is that these low-range images preserve edge information and highly distort luminance. Such spatial-comparison algorithms are scene dependent.³² Scene-dependent spatial processing found in painting, photography, and image processing are very successful at rendering HDR scenes.³ These spatial mechanisms render HDR images on devices with smaller dynamic range by preserving the scene's edge information and appearance. While scenes can have very large dynamic ranges, technology limits the range available in the display of these scenes. In evaluating the effectiveness of displays of HDR scenes, we need to evaluate the range visible to humans. This range is controlled by the interplay between the contrast and glare.

7 Conclusions

We have studied the HVR response to test HVS SD (0–2.7 log units) and DD (0–5.4 log units) targets with minimal changes in glare and simultaneous contrast. We studied targets with 100%, 50%, 8%, and 0% white backgrounds. Observers estimated appearances that were almost the same in both SD and DD displays with one exception. HDR images are limited by scene-dependent intraocular glare. In a white surround, with maximum glare, observers use an OD range of 2.0 log units to cover the range of appearances from white to black. By using half-white and half-black surrounds, we held simultaneous contrast constant and reduced the glare by one half. Observers use a range of 2.3 log units for white-to-black appearances. In a third experiment, we reduced the background to 8% white, decreasing glare further. Here, observers use a range of 2.9 log units for white to black. The exception was the 0% white background data. Observer data showed different tone-scale maps for SD and DD displays, even though their backgrounds appeared the same. They used a range of 2.7 log units for the SD target and a range of 5.0 log units for the DD target. Appearance in HDR images is controlled by both optical scattered light and spatial processing. Both mechanisms are scene-dependent. A single tone-scale function for luminance cannot describe the appearance that is controlled by two independent scene-dependent mechanisms: scatter and spatial processing. These mechanisms tend to cancel each other.

Acknowledgments

The authors wish to thank Marzia Pezzetti for his work in performing these measurements. We also want to thank Ivar Farup and Mary McCann for their help and discussions.

References

- 1 J. J. McCann, "Art, science, and appearance in HDR images," *J. Soc. Info. Display* **15/9**, 709–719 (2007).
- 2 J. J. McCann and A. Rizzi, "Camera and visual veiling glare in HDR images," *J. Soc. Info. Display* **15/9**, 721–730 (2007).
- 3 J. J. McCann, "Perceptual rendering of HDR in painting and photography," in *Proc. SPIE: Human Vision and Electronic Imaging XII*, B. Rogowitz, T. Pappas, and S. Daly, eds., 6806-30 (2008).
- 4 E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging Acquisition, Display and Image-Based Lighting* (Elsevier, Morgan Kaufmann, Amsterdam, 2006), Chap. 6.
- 5 P. E. Debevec and J. Malik, "Recovering high-dynamic range radiance maps from photographs," *ACM SIGGRAPH*, 369 (1997).
- 6 Laidler, Meiser, and Sanctuary (eds.), *Physical Chemistry* (Houghton Mifflin, Boston, 2003).
- 7 J. J. McCann and A. Rizzi, "Optical veiling glare limitations to in-camera scene radiance measurements," in: *ECVP 2006 Abstracts, Perception* **35**, Supplement, 51 (2006).
- 8 J. J. McCann and A. Rizzi, "Spatial Comparisons: The Antidote to Veiling Glare Limitations in HDR Images," in: *Proc ADEAC/SID & VESA*, 155–158 (2006).
- 9 J. J. McCann and A. Rizzi, "Veiling glare: the dynamic range limit of HDR images," in *SPIE Proc. Human Vision and Electronic Imaging XII*, B. Rogowitz, T. Pappas, and S. Daly (eds.), 6492-41 (2007).
- 10 J. J. McCann and A. Rizzi, "Spatial Comparisons: The Antidote to Veiling Glare Limitations in Image Capture and Display" in: *The Second Intl. Workshop on Image Media Quality and Its Applications*, E-1 (2007).
- 11 H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," *ACM Trans. Graphics* **23**(3), 760–768 (2004).
- 12 T. van den Berg, L. Van Rijn, R. Michael, C. Heine, T. Coeckelbergh, C. Nischler, H. Wilhelm, G. Grabner, M. Emesz, and R. Barraquer, "Straylight effects with aging and lens extraction," *Am. J. Ophthalmology* **144**, No. 3, 358–363 (2007).
- 13 A. Rizzi, M. Pezzetti, and J. J. McCann, "Separating the effects of glare from simultaneous contrast," in *SPIE Proc. Human Vision and Electronic Imaging XIII* 6806-09 (2008).
- 14 H. Davson, *The Eye: The Visual Process, Vol II* (Academic Press, New York, 1962).
- 15 A. Gilchrist, ed., *Lightness, Brightness and Transparency* (Lawrence Erlbaum Associates, Hillsdale, 1994).
- 16 J. J. McCann and R. L. Savoy, "Measurement of Lightness: Dependence on the position of a white in the field of view," in *Proc. SPIE Human Vision, Visual Processing and Digital Display II*, B. Rogowitz, ed., **1453**, 402–411 (1991).
- 17 J. J. McCann, "Visibility of gradients and low-spatial frequency sinusoids: Evidence for a distance constancy mechanism," *J. Photogr. Sci. Eng.* **22**, 64–68 (1978).
- 18 J. J. McCann and A. Rizzi, "The Spatial Properties of Contrast," in *Proc. IS&T/SID Color Imaging Conference*, **11**, 51–58 (2003).
- 19 H. W. Bodmann, P. Haubner, and A. M. Marsden, "A unified relationship between brightness and luminance," *CIE Proc.*, 99–102 (1979) and *Proc. CIE*, 50 (1979).
- 20 J. J. McCann, "Rendering high-dynamic range images: Algorithms that mimic human vision," in *Proc. AMOS Technical Conference*, 19–28 (2005).
- 21 J. J. McCann, "Aperture and Object Mode Appearances in Images," in *SPIE Proc.: Human Vision and Electronic Imaging XII*, B. Rogowitz, T. Pappas, and S. Daly, eds., 6292-26 (2007).
- 22 W. A. Stiehl, J. J. McCann, and R. L. Savoy, "Influence of intraocular scattered light on lightness-scaling experiments," *J. Opt. Soc. Am.* **73**, 1143–1148 (1983).
- 23 J. J. McCann and A. Rizzi, "Appearance of high-dynamic range images in a uniform lightness space," *Proc. CGIV 08/IS&T, 4th European Conference on Colour in Graphics* (2008).
- 24 A. Rizzi, M. Pezzetti, and J. J. McCann, "Separating the effects of glare from simultaneous contrast," in *SPIE Proc.: Human Vision and Electronic Imaging XIII*, **6806-09** (2008).
- 25 J. J. McCann, "Scene normalization mechanisms in humans," *Proc. IS&T/SID Second Color Imaging Conference*, 5–8 (1994).
- 26 J. J. McCann, "Lessons learned from Mondrians applied to real images and color gamuts," *Proc. IS&T/SID Seventh Color Imaging Conference*, 1–8 (1999).
- 27 A. Gilchrist, *Seeing Black and White* (Oxford University Press, Oxford, 2006).
- 28 J. J. Vos and T. van den Berg, CIE Research note 135/1, "Disability Glare," ISBN 3 900 734 97 6 (1999).
- 29 ISO 9358:1994 Standard, "Optics and optical instruments. Veiling glare of image forming systems. Definitions and methods of measurement" (ISO, 1994).

- 30 E. Land and J. J. McCann, "Lightness and Retinex theory," *J. Opt. Soc. Am.* **61**, 1–11 (1971).
- 31 T. Cornsweet and D. Teller, "Relation of increment thresholds to brightness and luminance," *J. Opt. Soc. Am.* **55**, 1303–1308 (1965).
- 32 J. J. McCann, ed., *J. Electron. Imaging* **13**, 1–145 (2004).

Alessandro Rizzi received a degree in computer science from the University of Milano and received his Ph.D. in information engineering from the University of Brescia (Italy). He taught information systems and computer graphics at the University of Brescia and at the Politecnico di Milano. Now, he is an assistant professor, teaching Multimedia and Human–Computer Interaction, and senior research fellow in the Department of Information Technologies at the University of Milano. Since 1990, he has been researching the field of digital imaging and vision. His main research area is the use of color information in digital images, with particular attention to color-perception mechanisms. He is the coordinator of the Italian Color Group.

John McCann received his B.A. degree in biology from Harvard University in 1964. He worked in, and later managed, the Vision Research Laboratory at Polaroid from 1961 to 1996. He has studied human color vision, digital image processing, large-format instant photography, and the reproduction of fine art. His 120 publications have studied Retinex theory, color from rod/Lcone interactions at low light levels, appearance with scattered light, and HDR imaging. He is a Fellow of IS&T. He is a past-president of IS&T and the Artists Foundation, Boston. He is currently consulting and continuing his research on color vision. In 1996, he received an SID Certificate of Commendation. He is the IS&T/OSA 2002 Edwin H. Land Medalist and IS&T 2005 Honorary Member, and in 2008 became Fellow of the Optical Society of America.