

A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes

Gregory Ward Larson, Holly Rushmeier, and Christine Piatko

Abstract—We present a tone reproduction operator that preserves visibility in high dynamic range scenes. Our method introduces a new histogram adjustment technique, based on the population of local adaptation luminances in a scene. To match subjective viewing experience, the method incorporates models for human contrast sensitivity, glare, spatial acuity, and color sensitivity. We compare our results to previous work and present examples of our techniques applied to lighting simulation and electronic photography.

Index Terms—Shading, image manipulation.

1 INTRODUCTION

THE real world exhibits a wide range of luminance values. The human visual system is capable of perceiving scenes spanning five orders of magnitude, and adapting more gradually to over nine orders of magnitude. Advanced techniques for producing synthetic images, such as radiosity and Monte Carlo ray tracing, compute the map of luminances that would reach an observer of a real scene. The media used to display these results—either a video display or a print on paper—cannot reproduce the computed luminances, or span more than a few orders of magnitude. However, the success of realistic image synthesis has shown that it is possible to produce images that convey the appearance of the simulated scene by mapping to a set of luminances that can be produced by the display medium. This is fundamentally possible because the human eye is sensitive to relative, rather than absolute, luminance values. However, a robust algorithm for converting real-world luminances to display luminances has yet to be developed.

The conversion from real-world to display luminances is known as *tone-mapping*. Tone-mapping ideas were originally developed for photography [1]. In photography (or video), chemistry (or electronics) are used, together with a human actively controlling the scene lighting and the camera, to map real-world luminances into an acceptable image on a display medium. In synthetic image generation, our goal is to avoid active control of lighting and camera settings. Furthermore, we hope to improve tone-mapping techniques by having direct numerical control over display values, rather than depending on the physical limitations of chemistry or electronics.

Consider a typical scene that poses a problem for tone reproduction in both photography and computer graphics image synthesis systems—a room illuminated by a window that looks out on a sunlit landscape. A human observer inside the room can easily see individual objects in the room, as well as features in the outdoor landscape. This is because the eye adapts locally as we scan the different regions of the scene. If we attempt to photograph our view, the result is disappointing. Either the window is overexposed and we can't see outside, or the interior of the room is underexposed and looks black. Current computer graphics tone operators either produce the same disappointing result, or introduce artifacts that do not match our perception of the actual scene.

In this paper, we present a new tone reproduction operator that reliably maps real-world luminances to display luminances, even in the problematic case just described. We consider the following two criteria most important for reliable tone-mapping:

- 1) Visibility is reproduced. You can see an object on the display if and only if you can see it in the real scene. Objects are not obscured in under- or overexposed regions, and features are not lost in the middle.
- 2) Viewing the image produces a subjective experience that corresponds with viewing the real scene. That is, the display should correlate well with memory of the actual scene. The overall impression of brightness, contrast, and color should be reproduced.

Previous tone-mapping operators have generally met one of these criteria at the expense of the other. For example, some preserve the visibility of objects while changing the impression of contrast, while others preserve the overall impression of brightness at the expense of visibility.

The new tone-mapping operator we present addresses our two criteria. We develop a method of modifying a luminance histogram, discovering clusters of adaptation levels and efficiently mapping them to display values to preserve local contrast visibility. We then use models for glare, color sensitivity, and visual acuity to reproduce imperfections in human vision that further affect visibility and appearance.

- G.W. Larson is with Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389. E-mail: gregl@sgi.com.
- H. Rushmeier is with the IBM T.J. Watson Research Center, 30 Saw Mill River Road, H4-C04, Hawthorne, NY 10532. E-mail: hertjwr@us.ibm.com.
- C. Piatko is with the Johns Hopkins University Applied Physics Laboratory, Johns Hopkins Road, Laurel, MD 20723-6099. E-mail: christine.piatko@jhuapl.edu.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number 105740.

2 PREVIOUS WORK

The high dynamic range problem was first encountered in computer graphics when physically accurate illumination methods were developed for image synthesis in the 1980s. (See Glassner [7] for a comprehensive review.) Previous methods for generating images were designed to automatically produce dimensionless values more or less evenly distributed in the range 0 to 1 or 0 to 255, which could be readily mapped to a display device. With the advent of radiosity and Monte Carlo path tracing techniques, we began to compute images in terms of real units with the real dynamic range of physical illumination. Fig. 1 is a false color image showing the magnitude and distribution of luminance values in a typical indoor scene containing a window to a sunlit exterior, as computed by the *Radiance* lighting simulation and rendering system [23]. The goal of image synthesis is to produce results, such as Fig. 4, which match our impression of what such a scene looks like. Initially though, researchers found that a wide range of displayable images could be obtained from the same input luminances—such as the unsatisfactory over- and underexposed linear reproductions of the image in Fig. 2 and Fig. 3.




Fig. 1. A false color image showing the world luminance values for a window office in candelas per meter squared (cd/m^2 or Nits).



Fig. 2. A linear mapping of the luminances in Fig. 1 that overexposes the view through the window.



Fig. 3. A linear mapping of the luminances in Fig. 1 that underexposes the view of the interior.



Fig. 4. The luminances in Fig. 1 mapped to preserve the visibility of both indoor and outdoor features using the new tone-mapping techniques described in this paper.

Initial attempts to find a consistent mapping from computed to displayable luminances were ad hoc and developed for computational convenience. One approach is to use a function that collapses the high dynamic range of luminance into a small numerical range. By taking the cube root of luminance, for example, the range of values is reduced to something that is easily mapped to the display range. This approach generally preserves visibility of objects, our first criterion for a tone-mapping operator. However, condensing the range of values in this way reduces fine detail visibility, and distorts impressions of brightness and contrast, so it does not fully match visibility or reproduce the subjective appearance required by our second criterion.

A more popular approach is to use an arbitrary linear scaling, either mapping the average of luminance in the real-world to the average of the display, or the maximum non-light-source luminance to the display maximum. For scenes with a dynamic range similar to the display device, this is successful. However, linear scaling methods do not maintain visibility in scenes with high dynamic range, since very bright and very dim values are clipped to fall within the display's limited dynamic range. Furthermore, scenes are mapped the same way, regardless of the absolute values

of luminance. A scene illuminated by a search light could be mapped to the same image as a scene illuminated by a flashlight, losing the overall impression of brightness and, so, losing the subjective correspondence between viewing the real and display-mapped scenes.

A tone-mapping operator proposed by Tumblin and Rushmeier [21] concentrated on the problem of preserving the viewer's overall impression of brightness. As the light level that the eye adapts to in a scene changes, the relationship between brightness (the subjective impression of the viewer) and luminance (the quantity of light in the visible range) also changes. Using a brightness function proposed by Stevens and Stevens [20], they developed an operator that would preserve the overall impression of brightness in the image, using one adaptation value for the real scene, and another adaptation value for the displayed image. Because a single adaptation level is used for the scene, though, preservation of brightness in this case is at the expense of visibility. Areas that are very bright or dim are clipped, and objects in these areas are obscured.

Ward [22] developed a simpler tone-mapping method, designed to preserve feature visibility. In this method, a nonarbitrary linear scaling factor is found that preserves the impression of contrast (i.e., the visible changes in luminance) between the real and displayed image at a particular fixation point. While visibility is maintained at this adaptation point, the linear scaling factor still results in the clipping of very high and very low values, and correct visibility is not maintained throughout the image.

Chiu et al. [2] addressed this problem of global visibility loss by scaling luminance values based on a spatial average of luminances in pixel neighborhoods. Values in bright or dark areas would not be clipped, but scaled according to different values based on their spatial location. Since the human eye is less sensitive to variations at low spatial frequencies than high ones, a variable scaling that changes slowly relative to image features is not immediately visible. However, in a room with a bright source and dark corners, the method inevitably produces display luminance gradients that are the opposite of real-world gradients. To make a dark region around a bright source, the transition from a dark area in the room to a bright area shows a decrease in brightness rather than an increase. This is illustrated in Fig. 5, which shows a bright source with a dark halo around it. The dark halo that facilitates rendering the visibility of the bulb disrupts what should be a symmetric pattern of light cast by the bulb on the wall behind it. The reverse gradient fails to preserve the subjective correspondence between the real room and the displayed image.

Inspired by the work of Chiu et al., Schlick [18] developed an alternative method that could compute a spatially varying tone-mapping. Schlick's work concentrated on improving computational efficiency and simplifying parameters, rather than improving the subjective correspondence of previous methods.

In the field of image processing, Jobson et al. [11] have developed digital tone-mapping methods for electronic photography based on Land's retinex theory [12]. The digital retinex techniques are similar in spirit to the method of Chiu et al., in that they effectively perform a local spatial



Fig. 5. Dynamic range compression based on a spatially varying scale factor (from [2]).

scaling. To avoid visible gradient reversals, the authors identify classes of images for which the techniques are effective, and empirically tune functions and parameters for these images. Retinex methods account well for color constancy effects (the independence of perceived object color from the spectral illumination of the object), but do not account for glare, acuity, or color sensitivity.

Contrast, brightness, and visibility are not the only perceptions that should be maintained by a tone-mapping operator. Nakamae et al. [16] and Spencer et al. [19] have proposed methods to simulate the effects of glare. These methods simulate the scattering in the eye by spreading the effects of a bright source in an image. Ferwerda et al. [5] proposed a method that accounts for changes in spatial acuity and color sensitivity as a function of light level. Our work is largely inspired by these papers, and we borrow heavily from Ferwerda et al. in particular. Besides maintaining visibility and the overall impression of brightness, we must include the effects of glare, spatial acuity, and color sensitivity to meet both our criteria for a successful operator.

A related set of methods for adjusting image contrast and visibility have been developed in the field of image processing for image enhancement (e.g., see Chapter 3 in [8]). Perhaps the best-known image enhancement technique is histogram equalization. In histogram equalization, the gray levels in an image are redistributed more evenly to make better use of the range of the display device. Numerous improvements have been made to simple equalization by incorporating models of perception. Frei [6] introduced histogram hyperbolization that attempts to redistribute perceived brightness, rather than screen gray levels. Frei approximated brightness using the logarithm of luminance. Subsequent researchers, such as Mokrane [14], have introduced methods that use more sophisticated models of perceived brightness and contrast.

The general idea of altering histogram distributions and using perceptual models to guide these alterations can be applied to tone-mapping. However, there are two important differences between techniques used in image enhancement and techniques for image synthesis and real-world tone-mapping:

- 1) In image enhancement, the problem is to correct an image that has already been distorted by photography or video recording and collapsed into a limited dynamic range. In our problem, we begin with an undistorted array of real-world luminances with a potentially high dynamic range.
- 2) In image enhancement, the goal is to take an imperfect image and *maximize* visibility or contrast. Maintaining subjective correspondence with the original view of the scene is irrelevant. In our problem, we want to maintain subjective correspondence. We want to *simulate* visibility and contrast, not maximize it. We want to produce visually accurate, not enhanced, images.

3 OVERVIEW OF THE NEW METHOD

In constructing a new method for tone-mapping, we wish to keep the elements of previous methods that have been successful and overcome the associated problems.

Consider again the room with a window looking out on a sunlit landscape. Like any high dynamic range scene, luminance levels occur in clusters, as shown in the histogram in Fig. 6, rather than being uniformly distributed throughout the dynamic range. The failure of any method that uses a single adaptation level is that it maps a large range of sparsely populated real-world luminance levels to a large range of display values. If the eye were sensitive to absolute values of luminance difference, this would be necessary. However, the eye is only sensitive to the fact that there are bright areas and dim areas. As long as the bright areas are displayed by higher luminances than the dim areas in the final image, the absolute value of the difference in luminance is not important. Exploiting this aspect of vision, we can close the gap between the display values for high- and low-luminance regions, and we have more display luminances to work with to render feature visibility.

Another failure of using a uniform adaptation level is that the eye rapidly adapts to the level of a relatively small angle in the visual field (i.e., about 1°) around the current fixation point [15]. When we look out the window, the eye adapts to the high exterior level, and, when we look inside, it adapts to the low interior level. Chiu et al. [2] attempted to account for this using spatially varying scaling factors, but this method can produce noticeable gradient reversals, as shown in Fig. 5.

Rather than adjusting the adaptation level based on spatial location in the image, we will base our mapping on the population of the luminance adaptation levels in the image. To identify clusters of luminance levels and *initially* map them to display values, we will use the cumulative distribution of the luminance histogram. More specifically, we will start with a cumulative distribution based on a logarithmic approximation of brightness from luminance values.

First, we calculate the population of levels from a luminance image of the scene, in which each pixel represents 1° in the visual field. By luminance, we specifically mean the measurable quantity that is the convolution of electromagnetic radiation with the standardized spectral sensitivity of the human eye, recorded in units of candelas/meter². (See Section 13.6.3 of [7] for a complete definition.) We want to

Histogram of Brightness

Window Office

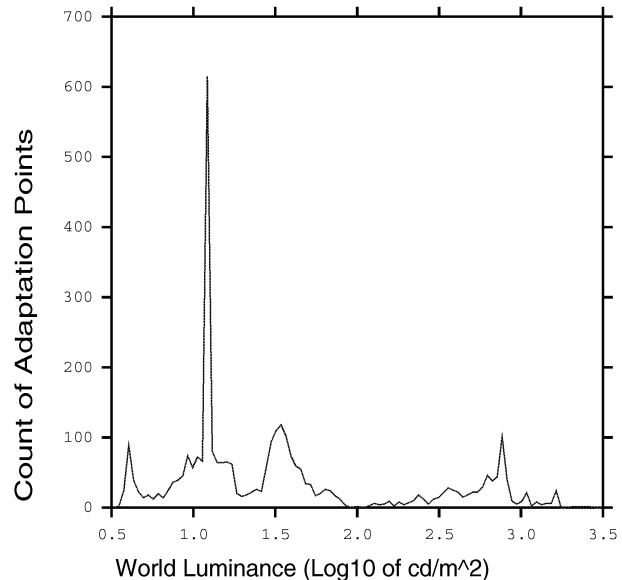


Fig. 6. A histogram of adaptation values from Fig. 1 (1° spot luminance averages).

compute a quantity to represent brightness, where brightness is the human subjective response to light which is not a physically measurable quantity. We make a crude approximation of the brightness values by taking the logarithm of luminance. (Note that we will not display logarithmic values, we will merely use them to obtain a distribution.) We then build a histogram and cumulative distribution function from these values. Since the brightness values are integrated over a small solid angle, they are, in some sense, based on a spatial average, and the resulting mapping will be local to a particular adaptation level. Unlike Chiu's method, however, the mapping for a particular luminance level will be consistent throughout the image, and will be order preserving. Specifically, an increase in real-scene luminance level will always be represented by an increase in display luminance. The histogram and cumulative distribution function will allow us to close the gaps of sparsely populated luminance values and avoid the clipping problems of single adaptation level methods. By deriving a single, global tone-mapping operator from locally averaged adaptation levels, we avoid the reverse gradient artifacts that can arise with a spatially varying multiplier.

We will use this histogram only as a starting point, and impose restrictions to preserve, rather than maximize, contrast based on models of human perception, using our knowledge of the true luminance values in the scene. Simulations of glare and variations in spatial acuity and color sensitivity will be added into the model to maintain subjective correspondence and visibility. In the end, we obtain a mapping of real world to display luminance similar to the one shown in Fig. 7.

World to Display Luminance Mapping

Window Office




Fig. 7. A plot comparing the global brightness mapping functions for Figs. 2, 3, and 4, respectively.

For our target display, all mapped brightness values below 1 cd/m^2 (0 on the vertical axis) or above 100 (2 on the vertical axis) are lost, because they are outside the displayable range. Here, we see that the dynamic range between 1.75 and 2.5 has been compressed, yet we don't notice it in the displayed result (Fig. 4). Compared to the two linear operators, our new tone-mapping is the only one that can represent the entire scene without losing object or detail visibility.

In the following section, we illustrate this technique for histogram adjustment based on contrast sensitivity. After this, we describe models of glare, color sensitivity, and visual acuity that complete our simulation of the measurable and subjective responses of human vision. Finally, we complete the methods presentation with a summary describing how all the pieces fit together.

4 HISTOGRAM ADJUSTMENT

In this section, we present a detailed description of our basic tone-mapping operator. We begin with the introduction of symbols and definitions, and a description of the histogram calculation. We then describe a naive equalization step that partially accomplishes our goals, but results in undesirable artifacts. This method is then refined with a linear contrast ceiling, which is further refined using human contrast sensitivity data.

4.1 Symbols and Definitions

- L_w = world luminance (in cd/m^2)
- B_w = world brightness, $\log(L_w)$
- L_{wmin} = minimum world luminance for scene
- L_{wmax} = maximum world luminance for scene
- L_d = display luminance (in cd/m^2)
- L_{dmin} = minimum display luminance (black level)
- L_{dmax} = maximum display luminance (white level)
- B_{de} = computed display brightness, $\log(L_d)$ (4)

- N = the number of histogram bins
- T = the total number of adaptation samples
- $f(b_i)$ = frequency count for the histogram bin at b_i
- Δb = the bin step size in $\log(\text{cd/m}^2)$
- $P(b)$ = the cumulative distribution function (2)
- $\log(x)$ = natural logarithm of x
- $\log_{10}(x)$ = decimal logarithm of x

4.2 Histogram Calculation

Since we are interested in optimizing the mapping between world adaptation and display adaptation, we start with a histogram of world adaptation luminances. The eye adapts for the best view in the fovea, so we compute each luminance over a 1° diameter solid angle corresponding to a potential foveal fixation point in the scene. We use a logarithmic scale for the histogram to best capture luminance population and subjective response over a wide dynamic range. This requires setting a minimum value as well as a maximum, since the logarithm of zero is $-\infty$. For the minimum value, we use either the minimum 1° spot average or 10^{-4} cd/m^2 (the lower threshold of human vision), whichever is larger. The maximum value is just the maximum spot average.

We start by filtering our original floating-point image down to a resolution that roughly corresponds to 1° square pixels. If we are using a linear perspective projection, the pixels on the perimeter will have slightly smaller diameter than the center pixels, but they will still be within the correct range. The following formula yields the correct resolution for 1° diameter pixels near the center of a linear perspective image:

$$S = 2 \tan(\theta/2) / 0.01745 \quad (1)$$

where

S = width or height in pixels

θ = horizontal or vertical full view angle

0.01745 = number of radians in 1°

For example, the view width and height for Fig. 4 are 63° and 45° , respectively, which yield a sample image resolution of 70 by 47 pixels. Near the center, the pixels will be 1° square exactly, but, near the corners, they will be closer to 0.85° for this wide-angle view. The filter kernel used for averaging will have little influence on our result, so long as every pixel in the original image is weighted similarly. We employ a simple box filter.

From our reduced image, we compute the logarithms of the floating-point luminance values. Here, we assume there is some method for obtaining the absolute luminances at each spot sample. If the image is uncalibrated, then the corrections for human vision will not work, although the method may still be used to optimize the visible dynamic range. (We will return to this in the summary.)

The histogram is taken between the minimum and maximum values mentioned earlier in equal-sized bins on a $\log(\text{luminance})$ scale. The algorithm is not sensitive to the number of bins, as long as there are enough to obtain adequate resolution. We use 100 bins in all of our examples. The resulting histogram for Fig. 1 is shown in Fig. 6.

4.2.1 Cumulative Distribution

The cumulative frequency distribution is defined as:

$$P(b) = \frac{\sum_{b_i < b} f(b_i)}{T} \quad (2)$$

where

$$T = \sum_{b_i} f(b_i)$$

(i.e., the total number of samples).

Later on, we will also need the derivative of this function. Since the cumulative distribution is a numerical integration of the histogram, the derivative is simply the histogram with an appropriate normalization factor. In our method, we approximate a continuous distribution and derivative by interpolating adjacent values linearly. The derivative of our function is:

$$\frac{dP(b)}{db} = \frac{f(b)}{T\Delta b} \quad (3)$$

where

$$\Delta b = \frac{[\log(L_{wmax}) - \log(L_{wmin})]}{N}$$

(i.e., the size of each bin).

4.3 Naive Histogram Equalization

If we wanted all the brightness values to have equal probability in our final displayed image, we could now perform a straightforward histogram equalization. Although this is not our goal, it is a good starting point for us. Based on the cumulative frequency distribution just described, the equalization formula can be stated in terms of brightness as follows:

$$B_{de} = \log(L_{dmin}) + [\log(L_{dmax}) - \log(L_{dmin})] \cdot P(B_w) \quad (4)$$

The problem with naive histogram equalization is that it not only compresses dynamic range (contrast) in regions where there are few samples, it also *expands* contrast in highly populated regions of the histogram. The net effect is to exaggerate contrast in large areas of the displayed image. Take, as an example, the scene shown in Fig. 8, with luminances computed using *Radiance*. Although we cannot see the region surrounding the lamps due to the clamped linear tone-mapping operator, the image appears to us as more or less normal. Applying the naive histogram equalization, Fig. 9 is produced. The tiles in the shower now have a mottled appearance. Because this region of world luminance values is so well represented, naive histogram equalization spreads it out over a relatively larger portion of the display's dynamic range, generating superlinear contrast in this region.

4.4 Histogram Adjustment With a Linear Ceiling

If the contrast being produced is too high, then what is an appropriate contrast for representing image features? The crude answer is that the contrast in any given region should not exceed that produced by a linear tone-mapping operator, since linear operators produce satisfactory results for scenes with limited dynamic range. We will take this simple approach first, and later refine our answer based on human contrast sensitivity.



Fig. 8. Rendering of a bathroom model mapped with a linear operator.



Fig. 9. Naive histogram equalization allows us to see the area around the light sources, but contrast is exaggerated in other areas, such as the shower tiles.

A linear ceiling on the contrast produced by our tone-mapping operator can be written thus:

$$\frac{dL_d}{dL_w} \leq \frac{L_d}{L_w} \tag{5a}$$

That is, the derivative of the display luminance with respect to the world luminance must not exceed the display luminance divided by the world luminance. Since we have an expression for the display luminance as a function of world luminance for our naive histogram equalization, we can differentiate the exponentiation of (4) using the chain rule and the derivative from (3) to get the following inequality:

$$\exp(B_{de}) \cdot \frac{f(B_w)}{T\Delta b} \cdot \frac{\log(L_{dmax}) - \log(L_{dmin})}{L_w} \leq \frac{L_d}{L_w} \tag{5b}$$

Since L_d is equal to $\exp(B_{de})$, this reduces to a constant ceiling on $f(b)$:

$$f(b) \leq \frac{T\Delta b}{\log(L_{dmax}) - \log(L_{dmin})} \tag{5c}$$

In other words, as long as we make sure no frequency count exceeds this ceiling, our resulting histogram will not exaggerate contrast.

How can we create this modified histogram? We considered both truncating larger counts to this ceiling and redistributing counts that exceeded the ceiling to other histogram bins. After trying both methods, we found truncation to be the simplest and most reliable approach. The only complication introduced by this technique is that once frequency counts are truncated, T changes, which changes the ceiling. We therefore apply iteration until a tolerance criterion is met, which says that fewer than 2.5 percent of the original samples exceed the ceiling.¹ Our pseudocode for `histogram_ceiling` is given below:

```

boolean function histogram_ceiling()
tolerance := 2.5% of histogram total
repeat {
  trimmings := 0
  compute the new histogram total T
  if T < tolerance then
    return FALSE
  foreach histogram bin i do
    compute the ceiling
    if  $f(b_i) > \text{ceiling}$  then {
      trimmings +=  $f(b_i) - \text{ceiling}$ 
       $f(b_i) := \text{ceiling}$ 
    }
  } until trimmings <= tolerance
return TRUE

```

This iteration will fail to converge (and the function will return `FALSE`) if and only if the dynamic range of the output device is already ample for representing the sample luminances in the original histogram. This is evident from (5c), since Δb is the world brightness range over the number of bins:

$$f(b_i) \leq \frac{T}{N} \cdot \frac{[\log(L_{wmax}) - \log(L_{wmin})]}{[\log(L_{dmax}) - \log(L_{dmin})]} \tag{5d}$$

If the ratio of the world brightness range over the display brightness range is less than one (i.e., our world range

fits in our display range), then our frequency ceiling is less than the total count over the number of bins. Such a condition will never be met, since a uniform distribution of samples would still be over the ceiling in every bin. It is easiest to detect this case at the outset, by checking the respective brightness ranges, and applying a simple linear operator if compression is unnecessary.



Fig. 10. Histogram adjustment with a linear ceiling on contrast preserves both lamp visibility and tile appearance.

Brightness Mapping Function Bathroom




Fig. 11. A comparison of naive histogram equalization (dashed line labeled “equalized”) with histogram adjustment (dotted line labeled “eq.linceil”). The linear mapping of brightness (solid line labeled “linear”) is also shown.

1. The tolerance of 2.5 percent was chosen as an arbitrary small value, and it seems to make little difference either to the convergence time or the results.

TABLE 1
PIECEWISE APPROXIMATION FOR $\Delta L_t(L_a)$

| log10 of just noticeable difference | applicable luminance range |
|--|---------------------------------------|
| -2.86 | $\log_{10}(L_a) < -3.94$ |
| $(0.405 \log_{10}(L_a) + 1.6)^{2.18} - 2.86$ | $-3.94 \leq \log_{10}(L_a) < -1.44$ |
| $\log_{10}(L_a) - 0.395$ | $-1.44 \leq \log_{10}(L_a) < -0.0184$ |
| $(0.249 \log(L_a) + 0.65)^{2.7} - 0.72$ | $-0.0184 \leq \log_{10}(L_a) < 1.9$ |
| $\log_{10}(L_a) - 1.255$ | $\log_{10}(L_a) \geq 1.9$ |

Once we have computed our modified histogram, the brightness mapping is obtained by substituting it back into (4). We call this method *histogram adjustment*, rather than histogram equalization, because the final brightness distribution is *not equalized*. The net result is a mapping of the scene's high dynamic range to the display's smaller dynamic range that minimizes visible contrast distortions, by compressing underrepresented regions without expanding overrepresented ones.

Fig. 10 shows the results of our histogram adjustment algorithm with a linear ceiling. The problems of exaggerated contrast are resolved, and we can still see the full range of brightness. A comparison of these tone-mapping operators is shown in Fig. 11. The naive operator is super-linear over a large range, seen as a very steep slope near world luminances around $10^{0.8}$.

The method we have just presented is itself quite useful. We have managed to overcome limitations in the dynamic range of typical displays without introducing objectionable contrast compression artifacts in our image. In situations where we want to get a good, natural-looking image without regard to how well a human observer would be able to see in a real environment, this may be an optimal solution. However, if we are concerned with reproducing both visibility and subjective experience in our displayed image, then we must take it a step further and consider the *limitations* of human vision.

4.5 Histogram Adjustment Based on Human Contrast Sensitivity

Although the human eye is capable of adapting over a very wide dynamic range, on the order of $1:10^9$, we do not see equally well at all light levels. As the light grows dim, we have more and more trouble detecting contrast. The relationship between adaptation luminance and the minimum detectable luminance change is well studied [3]. For consistency with earlier work, we use the same detection threshold function used by Ferwerda et al. [5]. This function covers sensitivity from the lower limit of human vision to daylight levels, and accounts for both rod and cone response functions. The piecewise fit is reprinted in Table 1.

We name this combined sensitivity function:

$$\Delta L_t(L_a) =$$

“just noticeable difference” for adaptation level L_a (6)

Ferwerda et al. did not combine the rod and cone sensitivity functions in this manner, since they used the two ranges for different tone-mapping operators. Since we are using this function to control the maximum reproduced contrast, we combine them at their crossover point of $10^{-0.0184} \text{ cd/m}^2$.

To guarantee that our display representation does not exhibit contrast that is more noticeable than it would be in the actual scene, we constrain the slope of our operator to the ratio of the two adaptation thresholds for the display and world, respectively. This is the same technique introduced by Ward [22] and used by Ferwerda et al. [5] to derive a global scale factor. In our case, however, the overall tone-mapping operator will not be linear, since the constraint will be met at all potential adaptation levels, not just a single selected one. The new ceiling can be written as:

$$\frac{dL_d}{dL_w} \leq \frac{\Delta L_t(L_d)}{\Delta L_t(L_w)} \quad (7a)$$

As before, we compute the derivative of the histogram equalization function, (4), to get:

$$\exp(B_{de}) \cdot \frac{f(B_w)}{T\Delta b} \cdot \frac{\log(L_{dmax}) - \log(L_{dmin})}{L_w} \leq \frac{\Delta L_t(L_d)}{\Delta L_t(L_w)} \quad (7b)$$

However, this time the constraint does not reduce to a constant ceiling for $f(b)$. We notice that, since L_d equals $\exp(B_{de})$ and B_{de} is a function of L_w from (4), our ceiling is completely defined for a given $P(b)$ and world luminance, L_w :

$$f(B_w) \leq \frac{\Delta L_t(L_d)}{\Delta L_t(L_w)} \cdot \frac{T\Delta b L_w}{[\log(L_{dmax}) - \log(L_{dmin})]} L_d \quad (7c)$$

where

$$L_d = \exp(B_{de}), B_{de} \text{ given in (4)}$$

Once again, we must iterate to a solution, since truncating bin counts will affect T and $P(b)$. We reuse the `histogram_ceiling` procedure given earlier, replacing the linear contrast ceiling computation with the above formula.

Fig. 12 shows the same curves for the linear tone-mapping and histogram adjustment with linear clamping shown before, in Fig. 11, but with the curve for naive histogram equalization replaced by our human visibility matching algorithm. We see the two histogram adjustment curves are very close. In fact, we would have some difficulty differentiating images mapped with our latest method and histogram adjustment with a linear ceiling. This is because the scene we have chosen has most of its luminance levels in the same range as our display luminances. Therefore, the ratio between display and world luminance detection thresholds is close to the ratio of the display and world adaptation luminances. This is known as Weber's law [25], and it holds true over a wide range of luminances where the eye sees equally well. This correspondence makes the right-hand sides of (5b) and (7b) equivalent, and so we should expect the same result as a linear ceiling.

Brightness Mapping Function

Bathroom




Fig. 12. Our tone-mapping operator, based on human contrast sensitivity (dashed line labeled “eq.hsens”), compared to the histogram adjustment with linear ceiling (dotted line labeld “eq.linceil”) used in Fig. 10. Human contrast sensitivity makes little difference at these light levels. The simple linear mapping is also shown here (solid line).

Brightness Mapping Function

Dim Bathroom




Fig. 13. The brightness map for the bathroom scene with lights dimmed to 1/100th of their original intensity, where human contrast sensitivity makes a difference. This difference is evident in the comparison of the linear map (dotted line labeled “eq.linceil”) and the human contrast sensitivity map (dashed line labeled “eq.hsens”). Again, the simple linear mapping is shown as a solid line for reference.

To see a contrast sensitivity effect, our world adaptation would have to be very different from our display adaptation. If we reduce the light level in the bathroom by a factor of 100, our ability to detect contrast is diminished. This shows up in a relatively larger detection threshold in the denominator of (7c), which reduces the ceiling for the frequency counts. The change in the tone-mapping operator is plotted in Fig. 13 and the resulting image is shown in Fig. 14.

Fig. 13 shows that the linear mapping is unaffected, since we just raise the scale factor to achieve an average exposure. Likewise, the histogram adjustment with a linear ceiling maps the image to the same display range, since its goal is to reproduce linear contrast. However, the ceiling based on human threshold visibility limits contrast over much of the scene, and the resulting image is darker and less visible everywhere except the top of the range, which is actually shown with higher contrast since we now have display range to spare.

Fig. 14 is darker and the display contrast is reduced compared to Fig. 10. Because the tone-mapping is based on *local adaptation* rather than a single global or spot average, threshold visibility is reproduced *everywhere* in the image, not just around a certain set of values. This criterion is met within the limitations of the display’s dynamic range.

5 HUMAN VISUAL LIMITATIONS

We have seen how histogram adjustment matches display contrast visibility to world visibility, but we have ignored three important limitations in human vision: glare, color sensitivity, and visual acuity. Glare is caused by bright sources in the visual periphery, which scatter light in the lens of the eye, obscuring foveal vision. Color sensitivity is reduced in dark environments, as the light-sensitive rods



Fig. 14. The dimmed bathroom scene mapped with the function shown in Fig. 13.

take over for the color-sensitive cone system. Visual acuity, the ability to resolve spatial detail, is also impaired in dark environments, due to the complete loss of cone response and the quantum nature of light sensation.

In our treatment, we will rely heavily on previous work performed by Moon and Spencer [15] and Ferwerda et al. [5], applying it in the context of a locally adapted visibility-matching model.

5.1 Veiling Luminance

Bright *glare sources* in the periphery reduce contrast visibility because light scattered in the lens obscures the fovea; this effect is less noticeable when looking directly at a source, since the eye adapts to the high light level. The influence of glare sources on contrast sensitivity is well studied and documented. We apply the original work of Holladay [10] and Moon and Spencer [15], which relates the effective adaptation luminance to the foveal average and glare source position and illuminance. A more precise model of veiling glare is offered by Spencer et al. [19], but the added computational expense is considerable.

In our presentation, we will first compute a low resolution “veil image” from our foveal sample values. The veil image represents luminance that has scattered within the eye. We will then interpolate this veil image to add glare effects to the original rendering. Finally, we will apply this veil as a correction to the adaptation luminances used for our contrast, color sensitivity, and acuity models.

Moon and Spencer base their formula for adaptation luminance on the effect of individual glare sources measured by Holladay, which they converted to an integral over the entire visual periphery. The resulting glare formula gives the effective adaptation luminance at a particular fixation for an arbitrary visual field:

$$L_a = 0.913L_f + \frac{K}{\pi} \iint_{\theta > \theta_f} \frac{L(\theta, \phi)}{\theta^2} \cos(\theta) \sin(\theta) d\theta d\phi \quad (8)$$

where

- L_a = corrected adaptation luminance (in cd/m^2)
- L_f = the average foveal luminance (in cd/m^2)
- $L(\theta, \phi)$ = the luminance in the direction (θ, ϕ)
- θ_f = foveal half angle, approx. 0.00873 radians (0.5°)
- K = constant measured by Holladay, 0.0096

The constant 0.913 in this formula is the remainder from integrating the second part, assuming one luminance everywhere. In other words, the periphery contributes less than 9 percent to the average adaptation luminance, due to the small value Holladay determined for K . If there are no bright sources, this influence can be safely neglected. However, bright sources will significantly affect the adaptation luminance, and should be considered in our model of contrast sensitivity.

To compute the veiling luminance corresponding to a given foveal sample (i.e., fixation point), we can convert the integral in (8) to an average over peripheral sample values:

$$L_{vi} = 0.087 \cdot \frac{\sum_{j \neq i} \frac{L_j \cos(\theta_{i,j})}{\theta_{i,j}^2}}{\sum_{j \neq i} \frac{\cos(\theta_{i,j})}{\theta_{i,j}^2}} \quad (9)$$

where

- L_{vi} = veiling luminance for fixation point i
- L_j = foveal luminance for fixation point j
- $\theta_{i,j}$ = angle between sample i and j (in radians)

Since we must compute this sum over all foveal samples j for each fixation point i , the calculation can be very time consuming. We minimize our costs by approximating the weight expression as:

$$\frac{\cos \theta}{\theta^2} \approx \frac{\cos \theta}{2 - 2\cos \theta} \quad (10)$$

Since the angles between our samples are most conveniently available as vector dot products, which is the cosine, the above weight computation is quite fast. However, for large images in terms of angular size, the L_{vi} calculation is still the most computationally expensive step in our method due to the double iteration over i and j .

To simulate the effect of glare on visibility, we simply add the computed veil map to our original image. Just as it occurs in the eye, the veiling luminance will obscure the visible contrast on the display by adding to both the background and the foreground luminance.² This was the original suggestion made by Holladay, who noted that the effect glare has on luminance threshold visibility is equivalent to what one would get by adding the veiling luminance function to the original image [10]. This is quite straightforward once we have computed our foveal-sampled veiling image given in (9). At each image pixel, we perform the following calculation:

$$L_{pvk} = 0.913L_{pk} + L_v(k) \quad (11)$$

where

- L_{pvk} = veiled pixel at image position k
- L_{pk} = original pixel at image position k
- $L_v(k)$ = interpolated veiling luminance at k

The $L_v(k)$ function is a simple bilinear interpolation on the four closest samples in our veil image computed in (9). The final image will be lighter around glare sources and just slightly darker *on* glare sources, since the veil is effectively being spread away from bright points. Although we have shown this as a luminance calculation, we retain color information, so that our veil has the same color cast as the responsible glare source(s).

Fig. 15 shows our original, fully lit bathroom scene again, this time adding in the computed veiling luminance. Contrast visibility is reduced around the lamps, but the veil falls off rapidly over other parts of the image. If we were to measure the luminance detection threshold at any given image point, the result should correspond closely to the threshold we would measure at that point in the actual scene. Due to the contrast compression necessary to fit this image within the dynamic range of the display, the subjective appearance of veil when looking at the light sources is incorrect. Ideally, we would adjust the display dynamically, based on the viewer's gaze, which would eliminate such artifacts.

2. The contrast is defined as the ratio of the foreground minus the background over the background, so adding luminance to both foreground and background reduces contrast.



Fig. 15. Our tone reproduction operator for the original bathroom scene with veiling luminance added.



Fig. 16. Our dimmed bathroom scene with tone-mapping using human contrast sensitivity, veiling luminance, and mesopic color response.

Although we can reproduce visibility with this method, we cannot reproduce the physical discomfort associated with real glare situations, and, without it, the subjective correspondence is lacking. We cannot overcome this limitation with conventional display methods, because conventional displays cannot reproduce the sometimes painful luminance differences present in real scenes.

Since glare sources scatter light onto the fovea, they also affect the local adaptation level, and we should consider this in the other parts of our calculation. We therefore apply the computed veiling luminances to our foveal samples as a correction *before* the histogram generation and adjustment described in Section 4. We deferred the introduction of this correction factor to simplify our presentation, since, in most cases, it only weakly affects the brightness mapping function.

The correction to local adaptation is the same as (11), but without interpolation, since our veil samples correspond one-to-one:

$$L_{ai} = 0.913L_i + L_{vi} \tag{12}$$

where

- L_{ai} = adjusted adaptation luminance at fixation point i
- L_i = foveal luminance for fixation point i

We will also employ these L_{ai} adaptation samples for the models of color sensitivity and visual acuity that follow.

5.2 Color Sensitivity

To simulate the loss of color vision in dark environments, we use the technique presented by Ferwerda et al. [5] and ramp between a scotopic (gray) response function and a

photopic (color) response function as we move through the mesopic range. The lower limit of the mesopic range, where cones are just starting to get enough light, is approximately 0.0056 cd/m^2 . Below this value, we use the straight scotopic luminance. The upper limit of the mesopic range, where rods are no longer contributing significantly to vision, is approximately 5.6 cd/m^2 . Above this value, we use the straight photopic luminance plus color. In between these two world luminances (i.e., within the mesopic range), our adjusted pixel is a simple interpolation of the two computed output colors, using a linear ramp based on luminance.

Since we do not have a value available for the scotopic luminance at each pixel, we used a least squares fit to the colors on the Macbeth ColorChecker Chart™. (See [7] for the appropriate spectral curves.) The approximate relation is given below:

$$Y_{scot} \approx Y \cdot \left[1.33 \cdot \left(1 + \frac{Y + Z}{X} \right) - 1.68 \right] \tag{13}$$

where

- Y_{scot} = scotopic luminance
- X, Y, Z = photopic color, CIE 2° observer (Y is luminance)

This is a very good approximation to scotopic luminance for most natural colors, and it avoids the need to render another channel. We also have an approximation based on RGB values, but, since there is no accepted standard for RGB primaries in computer graphics, this is much less reliable.

Fig. 16 shows our dimmed bathroom scene with the human color sensitivity function in place. Notice there is still some veiling, even with the lights reduced to 1/100th

their normal level. This is because the relative luminances are still the same, and they scatter in the eye as before. The only difference here is that the eye cannot adapt as well when there is so little light, so everything appears dimmer, including the lamps. The colors are clearly visible near the light sources, but gradually less visible in the darker regions.

5.3 Visual Acuity

Besides losing the ability to see contrast and color, the human eye loses its ability to resolve fine detail in dark environments. The relationship between adaptation level and foveal acuity has been measured in subject studies reported by Shaler [26]. At daylight levels, human visual acuity is very high, about 50 cycles/degree. In the mesopic range, acuity falls off rapidly from 42 cycles/degree at the top down to four cycles/degree near the bottom. Near the limits of vision, the visual acuity is only about two cycles/degree. Shaler's original data is shown in Fig. 17, along with the following functional fit:

$$R(L_a) \approx 17.25 \arctan(1.4 \log_{10}(L_a) + 0.35) + 25.72 \quad (15)$$

where

$R(L_a)$ = visual acuity in cycles/degree

L_a = local adaptation luminance (in cd/m^2).

Human Visual Acuity Function (Foveal) due to Shaler




Fig. 17. Shaler's visual acuity data and our functional fit to it.

In their tone-mapping paper, Ferwerda et al. applied a global blurring function based on a single adaptation level [5]. Since we wish to adjust for acuity changes over a wide dynamic range, we must apply our blurring function locally, according to the foveal adaptation computed in (12). To do this, we implement a variable-resolution filter, using an image pyramid and interpolation, which is the *mip map* introduced by Williams [24] for texture mapping. The only difference here is that we are working with real values rather than integers.

At each point in the image, we interpolate the local acuity based on the four closest (veiled) foveal samples and Shaler's data. It is very important to use the foveal data (L_{ai}) and not the original pixel value, since it is the fovea's adaptation that determines acuity. The resulting image will show higher resolution in brighter areas, and lower resolution in darker areas.

Fig. 18 shows our dim bathroom scene again, this time applying the variable acuity operator applied together with all the rest. Since the resolution of the printed image is low, we enlarged two areas for a closer look. The bright area has an average level around $25 \text{ cd}/\text{m}^2$, corresponding to a visual acuity of about 45 cycles/degree. The dark area has an average level of around $0.05 \text{ cd}/\text{m}^2$, corresponding to a visual acuity of about nine cycles/degree. Unlike the results shown in Fig. 18, the global averaging of Ferwerda et al. [5] would have resulted in the same degree of blurring in both regions.




Fig. 18. The dim bathroom scene with variable acuity adjustment. The insets show two areas, one light and one dark, and the relative blurring of the two.

6 METHOD SUMMARY

We have presented a method for matching the visibility of high dynamic range scenes on conventional displays, accounting for human contrast sensitivity, veiling luminance, color sensitivity, and visual acuity, all in the context of a local adaptation model. However, in presenting this method in parts, we have not given a clear idea of how the parts are integrated together into a working program.

The order in which the different processes are executed to produce the final image is import. These are the steps in the order they are usually performed:

```
procedure match_visibility()
  compute 1° foveal sample image
  compute veil image
```

```

add veil to foveal adaptation image
add veil to image
blur image locally based on visual acuity function
apply color sensitivity function to image
generate histogram of effective adaptation image
adjust histogram to contrast sensitivity function
apply histogram adjustment to image
translate CIE results to display RGB values
end

```

We have not discussed the final step, mapping the computed display luminances and chrominances to appropriate values for the display device (e.g., monitor RGB settings). This is a well-studied problem, and we refer the reader to the literature (e.g., [9]) for details. Bear in mind that the mapped image accounts for the black level of the display, which must be subtracted out before applying the appropriate gamma and color corrections.

A few of the steps in this sequence may be moved around, or removed entirely for a different effect. Specifically, it makes little difference whether the luminance veil is added before or after the blurring function, since the veil varies slowly over the image. Also, the color sensitivity function may be applied anywhere after the veil is added, so long as it is before histogram adjustment.

If the goal is to optimize visibility and appearance without regard to the limitations of human vision, then all the steps between computing the foveal average and generating the histogram may be skipped, and a linear ceiling may be applied during histogram adjustment, instead of the human contrast sensitivity function. The result will be an image with all parts visible on the display, regardless of the world luminance level or the presence of glare sources. If the goal is just to produce a reasonable image when the absolute luminance levels and spectral distributions are unknown, the histogram can be formed from gray level values (i.e., from simple weighted averages of red, green, and blue).

7 RESULTS

In our dynamic range compression algorithm, we have exploited the fact that humans are insensitive to relative and absolute differences in luminance. For example, we can see that it is brighter outside than inside on a sunny day, but we cannot tell how much brighter (three times or 100) or what the actual luminances are (10 cd/m^2 or 10,000). With the additional display range made available by adjusting the histogram to close the gaps between luminance levels, visibility (i.e., contrast) within each level can be properly preserved. Furthermore, this is done in a way that is compatible with subjective aspects of vision.

In the development sections, two synthetic scenes have served as examples. In this section, we show results from two different application areas—lighting simulation using *Radiance*, and electronic photography.

7.1 Lighting Simulation

In lighting design, it is important to simulate what it is like to *be* in an environment, not what a photograph of the environment looks like. Figs. 19 and 20 show examples of real lighting design applications.



Fig. 19. A simulation of a shipboard control panel under emergency lighting.




Fig. 20. A simulation of an air traffic control console.

In Fig. 19, the emergency lighting of a control panel is shown. It is critical that the lighting provide adequate visibility of signage and levers. An image synthesis method that cannot predict human visibility is useless for making lighting or system design judgments.

Fig. 20 shows a flight controller's console. Being able to switch back and forth between the console and the outdoor view is an essential part of the controller's job. Again, judgments on the design of the console cannot be made on the basis of ill-exposed or arbitrarily mapped images.

Fig. 21 is not a real lighting application, but represents another type of interesting lighting. In this case, the high dynamic range is not represented by large areas of either high or low luminance. Very high, almost point, luminances are scattered in the scene. The new tone-mapping works equally well on this type of lighting, preserving visibility while keeping the impression of the brightness of the point sources. The color sensitivity and variable acuity mapping also correctly represent the sharp color view of areas surrounding the lights, and the grayed blurring of more dimly lit areas.

Each of these images contains about a million pixels and took under 30 seconds to process on an SGI O2 R5000. This



Fig. 21. A Christmas tree with very small light sources.

represents a tiny fraction of the time required to compute the original images using *Radiance*.

7.2 Electronic Photography

Finally, we present an example from electronic photography. In traditional photography, it is impossible to set the exposure so all areas of a scene are visible as they would be to a human observer. New techniques of digital compositing are now capable for creating images with much higher dynamic ranges. Our tone reproduction operator can be applied to appropriately map these images into the range of a display device.

Fig. 22 shows the interior of a church, taken on print film by a 35mm SLR camera with a 15mm fisheye lens. The stained glass windows are not completely visible because the recording film has been saturated, even though the rafters on the right are too dark to see. Fig. 23 shows our tone reproduction operator applied to a high dynamic range version of this image, called a *radiance map*. The radiance map was generated from 16 separate exposures, each separated by one stop. These images were scanned, registered, and the full dynamic range was recovered using an algorithm developed by Debevec and Malik [4]. Our tone-mapping operator makes it possible to retain the image features shown in Fig. 23, whose world luminances span over six orders of magnitude.

Fig. 23 contains about 400,000 pixels, and took under five seconds to process on an SGI O2. We can process the



Fig. 22. A scanned photograph of Memorial Church.



Fig. 23. Histogram adjusted radiance map of Memorial Church.

same image, using integer arithmetic and lookup tables, in about one second.

The field of electronic photography is still in its infancy. Manufacturers are rapidly improving the dynamic range of sensors and other electronics that are available at a reasonable cost. Visibility preserving tone reproduction operators will be essential in accurately displaying the output of such sensors in print and on common video devices.

8 CONCLUSIONS AND FUTURE WORK

There are still several degrees of freedom possible in this tone-mapping operator. For example, the method of computing the foveal samples corresponding to viewer fixation points could be altered. This would depend on factors such as whether an interactive system or a preplanned animation is being designed. Even in a still image, a theory of where the observer is likely to focus attention could be applied to improve the initial adaptation histogram. Additional modifications could easily be made to the threshold sensitivity, veil, and acuity models to simulate the effects of aging and certain types of visual impairment.

This method could also be extended to other application areas. The tone-mapping could be incorporated into global illumination calculations to make them more efficient by relating error to visibility. The mapping could also become part of a metric to compare images and validate simulations, since the results correspond roughly to human perception [17].

Some of the approximations in our operator merit further study, such as color sensitivity changes in the mesopic range. A simple choice was made to interpolate linearly between scotopic and photopic response functions, which follows Ferwerda et al. [5], but should be examined more closely. The effect of the luminous surround on adaptation should also be considered, especially for projection systems in darkened rooms. Finally, the current method pays little attention to absolute color perception, which is strongly affected by global adaptation and source color (i.e., white balance).

The examples and results we have shown match well with the subjective impression of viewing the actual environments being simulated or recorded. On this informal level, our tone-mapping operator has been tested experimentally. To improve upon this, more rigorous validations are needed. While validations of image synthesis techniques have been performed before (e.g., Meyer et al. [13]), they have not dealt with the level of detail required for validating an accurate tone operator. Validation experiments will require building a stable, nontrivial, high dynamic range environment and introducing observers to the environment in a controlled way. Reliable, calibrated methods are needed to capture the actual radiances in the scene and reproduce them on a display following the tone-mapping process. Finally, a series of unbiased questions must be formulated to evaluate the subjective correspondence between observation of the physical scene and observation of images of the scene in various media. While such experiments will be a significant undertaking, the level of sophistication in image synthesis and electronic photography requires such detailed validation work.

The dynamic range of an interactive display system is limited by the technology required to control continual,

intense, focused energy over millisecond time frames, and by the uncontrollable elements in the ambient viewing environment. The technological, economic, and practical barriers to display improvement are formidable. Meanwhile, luminance simulation and acquisition systems continue to improve, providing images with higher dynamic range and greater content, and we need to communicate this content on conventional displays and hard copy. To encourage further experimentation in tone-mapping, the authors have developed a new, high dynamic range image format, and published more than 100 sample images at the following web site:

<http://www.sgi.com/Technology/pixformat/>

ACKNOWLEDGMENTS

The authors wish to thank Robert Clear, Samuel Berman, Charles Fenimore, Peter Shirley, and a host of anonymous reviewers (you know who you are) for their helpful input. This work was supported by the Laboratory Directed Research and Development Funds of Lawrence Berkeley National Laboratory under the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

REFERENCES

- [1] C.J. Bartelton and E.J. Breneman, "Brightness Reproduction in the Photographic Process," *Photographic Science and Eng.*, vol. 11, pp. 254-262, July-Aug. 1967.
- [2] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, "Spatially Nonuniform Scaling Functions for High Contrast Images," *Proc. Graphics Interface '93*, pp. 245-253, Toronto, Canada, May 1993.
- [3] CIE, *An Analytical Model for Describing the Influence of Lighting Parameters Upon Visual Performance*, vol. 1. Technical Foundations, CIE 19/2.1, Technical Committee 3.1, 1981.
- [4] P. Debevec and J. Malik, "Recovering High Dynamic Range Radiance Maps from Photographs," *Proc. ACM SIGGRAPH '97*.
- [5] J. Ferwerda, S. Pattanaik, P. Shirley, and D.P. Greenberg, "A Model of Visual Adaptation for Realistic Image Synthesis," *Proc. ACM SIGGRAPH '96*, p. 249-258, 1996.
- [6] W. Frei, "Image Enhancement by Histogram Hyperbolization," *Computer Graphics and Image Processing*, vol. 6, pp. 286-294, 1977.
- [7] A. Glassner, *Principles of Digital Image Synthesis*. San Francisco: Morgan Kaufman, 1995.
- [8] W. Green, *Digital Image Processing: A Systems Approach*. New York: Van Nostrand Reinhold, 1983.
- [9] R. Hall, *Illumination and Color in Computer Generated Imagery*. New York: Springer-Verlag, 1989.
- [10] L.L. Holladay, *J. Optical Soc. Am.*, vol. 12, p. 271, 1926.
- [11] D.J. Jobson, Z. Rahman, and G.A. Woodell, "Properties and Performance of a Center/Surround Retinex," *IEEE Trans. Image Processing*, vol. 6, no. 3, pp. 451-462, Mar. 1997.
- [12] E.H. Land, "The Retinex Theory of Color Vision," *Scientific American*, vol. 237, no. 6, 1977.
- [13] G. Meyer, H. Rushmeier, M. Cohen, D. Greenberg, and K. Torrance, "An Experimental Evaluation of Computer Graphics Imagery," *ACM Trans. Graphics*, vol. 5, no. 1, pp. 30-50, Jan. 1986.
- [14] A. Mokrane, "A New Image Contrast Enhancement Technique Based on a Contrast Discrimination Model," *CVGIP: Graphical Models and Image Processing*, vol. 54, no. 2, pp. 171-180, Mar. 1992.
- [15] P. Moon and D. Spencer, "The Visual Effect of Non-Uniform Surrounds," *J. Optical Soc. Am.*, vol. 35, no. 3, pp. 233-248, 1945.
- [16] E. Nakamae, K. Kaneda, T. Okamoto, and T. Nishita, "A Lighting Model Aiming at Drive Simulators," *Proc. ACM SIGGRAPH '90*, vol. 24, no. 3, pp. 395-404, June 1990.
- [17] H. Rushmeier, G. Ward, C. Piatko, P. Sanders, and B. Rust, "Comparing Real and Synthetic Images: Some Ideas about Metrics," *Proc. Sixth Eurographics Workshop Rendering*, Dublin, Ireland, June 1995.

- [18] C. Schlick, "Quantization Techniques for Visualization of High Dynamic Range Pictures," *Photorealistic Rendering Techniques*, G. Sakas, P. Shirley, and S. Mueller, eds., pp. 7-20. Berlin: Springer-Verlag, 1995.
- [19] G. Spencer, P. Shirley, K. Zimmerman, and D. Greenberg, "Physically-Based Glare Effects for Computer Generated Images," *Proc. ACM SIGGRAPH '95*, pp. 325-334, 1995.
- [20] S.S. Stevens and J.C. Stevens, "Brightness Function: Parametric Effects of Adaptation and Contrast," *J. Optical Soc. Am.*, vol. 53, p. 1,139, 1960.
- [21] J. Tumblin and H. Rushmeier. "Tone Reproduction for Realistic Images," *IEEE Computer Graphics and Applications*, vol. 13, no. 6, pp. 42-48, Nov. 1993.
- [22] G. Ward, "A Contrast-Based Scalefactor for Luminance Display," *Graphics Gems IV*, P.S. Heckbert, ed. Boston: Academic Press.
- [23] G.J. Ward, "The RADIANCE Lighting Simulation and Rendering System," *Computer Graphics*, pp. 459-472, July 1994.
- [24] L. Williams, "Pyramidal Parametrics," *Computer Graphics*, vol. 17, no. 3, July 1983.
- [25] L.A. Riggs, "Vision," *Woodworth and Schlosber's Experimental Psychology*, third ed., J.W. Kling and L.A. Riggs, eds. New York: Holt, Rinehart, and Winston, 1971.
- [26] Shaler, "The Relation Between Visual Acuity and Illimination," *J. General Psychology*, vol. 21, pp. 165-188, 1937.



Gregory Ward Larson has an AB in physics from the University of California at Berkeley, and an MS in computer science from San Francisco State University. He is a member of the technical staff at Silicon Graphics, Inc. His professional interests include physically based rendering, digital photography, electronic data standards, and energy conservation, and he has published numerous papers in computer graphics and illumination engineering. He is the primary author of the *Radiance* lighting simulation and rendering system, which he developed while at Lawrence Livermore Berkeley National Laboratories in Berkeley, and EPFL in Lausanne.



Holly Rushmeier received the BS, MS, and PhD degrees in mechanical engineering from Cornell University in 1977, 1986, and 1988, respectively. She is a research staff member at the IBM T.J. Watson Research Center. Since receiving the PhD, she has held positions at the Georgia Institute of Technology, and at the National Institute of Standards and Technology. In 1990, she was selected as a U.S. National Science Foundation Presidential Young Investigator. In 1996, she served as the papers chair for the ACM SIGGRAPH conference. She is currently editor-in-chief of *ACM Transactions on Graphics*. Her research interests include data visualization and synthetic image generation.



Christine Piatko received her BA degree in computer science and mathematics from New York University in 1986, and her MS and PhD degrees from Cornell University in 1989 and 1993, respectively. She is currently a senior computer science researcher at the Johns Hopkins University Applied Physics Laboratory, and previously worked at the National Institute of Standards and Technology. Her research interests include computational geometry and visualization.