Appearance-based Image Splitting for HDR Displays

Dan Zhang, James A. Ferwerda and Jonathan B. Phillips; Munsell Color Science Laboratory, Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, Rochester, NY 14623, USA

Abstract

High dynamic range (HDR) displays that incorporate two optically-coupled image planes have recently been developed. The existence of such displays requires HDR image splitting algorithms that send appropriate signals to each plane to faithfully reproduce the appearance of the HDR input. In this paper we introduce a new appearance-based HDR image splitting algorithm that incorporates the iCAM06 image appearance model to do image enhancement. We compare its performance to the widely used luminance square root algorithm and report the results of image quality experiments that compare the two algorithms with respect to contrast, color, sharpness, naturalness, and overall quality. We find that the overall quality, color, and naturalness of the images produced by the new algorithm is superior to those produced by the square root method. The new algorithm provides a principled and effective approach for presenting HDR images on dual imager HDR displays.

Introduction

Real-world scenes encompass a 10 log unit range of luminance levels, from below 0.001 cd/m^2 to over 100,000 cd/m^2 [1]. The 4 to 6 log unit (10,000 to 1,000,000 to 1) luminance dynamic range found in many scenes vastly exceeds the ranges that can be captured or reproduced by conventional imaging systems. For example, images captured by conventional digital cameras typically have dynamic ranges between 100 or 1000 to 1 around a level set by aperture and shutter speed. Conventional display systems are similar in that their output dynamic ranges are on the order of 100 to 1 with maximum luminance between 100 and 400 cd/m^2 [2].

Over the past fifteen years, new technologies have been developed for overcoming this dynamic range bottleneck to produce image capture and display systems capable of recording and reproducing high dynamic range (HDR) images. High dynamic range image capture systems have been developed for both still images [3,4,5,6,7] and video recordings [8]. For displaying HDR images, several alternative systems have been developed, including both softcopy [9,10] and hardcopy [11,12] designs.

HDR display systems typically reproduce high dynamic range images by using two standard dynamic range (SDR) imagers that are optically coupled. The basic principle is that one imager (such as a projector or LED array) provides spatially varying illumination for a second imager (for example a transmissive LCD or reflective print), allowing HDR image values to be reproduced.

This dual image plane design requires that a given HDR input image be split into two complementary SDR components that drive the coupled systems. The widely used *square root HDR splitting algorithm* first converts an input HDR image to XYZ tristimulus values, then takes the square root of the Y channel and sends this achromatic signal to one image plane. A color signal is created by composing $\sqrt{Y}$ with its corresponding X and Z channels to the other. Under ideal conditions, this approach will reproduce the original luminance range of the HDR input, but faithful color reproduction is not considered.

To take a more principled approach to the HDR image splitting problem, we have developed a new algorithm based on the iCAM06 image appearance model [13]. The algorithm first uses iCAM06 to create a SDR color image that is sent to one plane, then calculates a luminance residual that is sent to the other plane to reproduce the HDR luminance range [14]. The goal of the algorithm is to create displayed HDR images that better reproduce the visual appearances of HDR scenes.

In the following sections, we first describe prior work on HDR displays and the luminance square root HDR image splitting algorithm, we then describe our new appearance-based HDR image splitting algorithm, and report on a series of experiments designed to evaluate and compare the appearance of HDR images displayed using the two algorithms.

Related work

HDR display systems

HDR displays can be classified as softcopy and hardcopy. In 2004 Seetzen et al. [9] described two designs for softcopy HDR displays that trans-illuminate an image displayed on an LCD panel with a locally-dimmable backlight provided by a DLP projector or LED array. Visser et al. [10] described an alternate design that used two stacked LCD panels with a bright uniform backlight.

In terms of hardcopy HDR displays, Ledda et al. [11] introduced a HDR still image viewer that used binocular optics, layered transparencies and a uniform incandescent backlight. Bimber et al. [12] developed a print-based HDR display that used a video projector to superimpose spatially varying illumination on a reflective photo print.
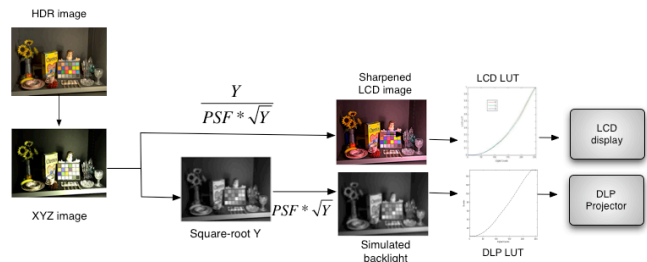

Figure 1: Luminance square-root HDR image splitting algorithm.

HDR image splitting algorithms

To drive a dual plane HDR display, the input HDR image must be split into the two SDR components that are sent to each plane. The explicit goal of HDR image splitting algorithms is to

accurately reproduce the input HDR image values on the display. The implicit goal is to reproduce the visual appearance of the scene. Work is just beginning on HDR image splitting algorithms and some of this work is proprietary, but three algorithms have been published.

The most widely used HDR image splitting method is the *luminance square root algorithm* [15]. A flowchart of the algorithm is shown in Figure 1. The algorithm starts with the assumption that the two planes in the display (shown in the Figure as a color LCD front plane and achromatic DLP projector backlight) have linear response functions and equal dynamic range. The input HDR image is first transformed to XYZ tristimulus values. Then the square root of the Y channel is calculated and two SDR images are created. One is an achromatic image that corresponds to the $\sqrt{Y}$ channel. The other is a color image that is formed by dividing the original Y channel by $\sqrt{Y}$ and composing the result with its corresponding X and Z channels. Because the backlight plane typically has lower spatial resolution than the front plane, the image sent to the front plane is spatially sharpened by the inverse of the point spread function (*PSF*) of the backlight. Finally, both images are converted to RGB and corrected by each display's response LUT. Assuming ideal display response properties, the square root method should be able to accurately reproduce the luminance values of any input image that falls within the HDR display's dynamic range. However, accurate color reproduction is not considered and loss of color saturation is a common artifact associated with this algorithm. Recently Luka and Ferwerda [16] introduced a variation on the square root algorithm that accommodates HDR displays with imagers that have unequal response properties. By transitioning from a square root function to a linear function at low luminances, they increased gamut utilization and improved the saturation of dark colors.

Guarnieri et al. [17,18] developed an HDR display for radiological applications by layering two high quality grayscale medical LCD displays. Due to the critical nature of the application, they were concerned with the accuracy of the displayed image and the effects of the image splitting algorithm on the visibility of image features. They developed an *optimization-based HDR image splitting algorithm* that simultaneously considered luminance reconstruction errors and spatial parallax errors caused by the thickness of the layered LCD panels. They approached HDR image splitting as an optimization problem and have produced an algorithm that typically achieves perfect luminance reconstruction and minimal parallax errors. When a perfect reconstruction is not possible, errors are minimized through the application of a visible difference metric. While this algorithm is based on sound mathematical and perceptual principles, it is designed to handle grayscale radiological images shown on a dual layer LCD display and is not directly applicable to the general class of color HDR images or other HDR display technologies.

Appearance-based HDR image splitting

The goal of this paper is to develop an HDR image splitting algorithm that produces displayed HDR images that better capture the appearance of HDR scenes. To achieve this goal, we have developed a new algorithm that is based on the iCAM06 image appearance model. The algorithm is illustrated in Figure 2 and described below.
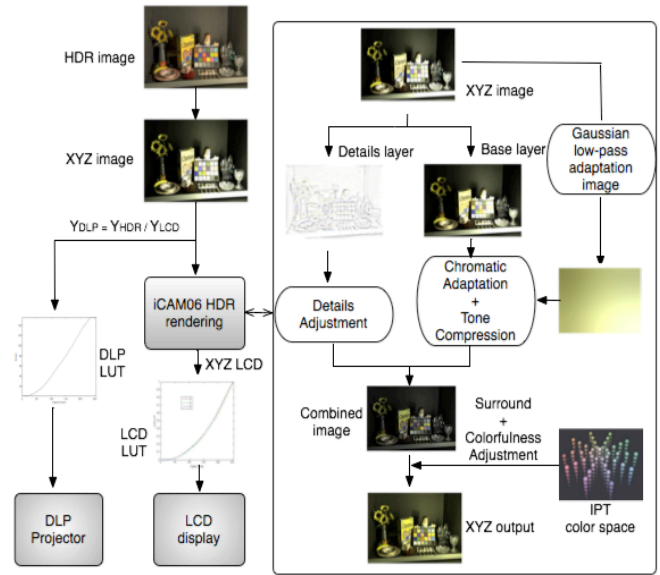

Figure 2: Appearance-based HDR image splitting algorithm.

First, the input HDR image is transformed to XYZ tristimulus values. Next the XYZ image is split into two streams. To create the signal for the LCD front plane, the XYZ image is processed through the iCAM06 HDR image rendering algorithm. Within this module, the image is split by a bilateral filter into base and details layers. A low pass version of the image is also calculated. This low pass image is then used to transform the base layer to account for the effects of luminance and color adaptation. The adapted base layer is then recombined with the details layer and this now SDR image is then transformed to the IPT opponent color space where the effects of the image surround are taken into account and colorfulness is adjusted. Finally the tone compressed, color adjusted SDR image is transformed back to RGB, corrected by the LCD's LUT, and sent to the display's front plane. In the other stream, the Y channel of the original XYZ image is split off, and a HDR luminance residual signal is calculated by dividing the the original HDR Y channel (Y_{HDR}) by the tone compressed Y channel (Y_{LCD}) of the image being sent to the LCD. This achromatic Y residual image (Y_{DLP}) is then converted to RGB and sent to the backlight imager (in this case a DLP projector).

The advantage of this algorithm is that the HDR image is mapped to the display using an image appearance model (iCAM06) that accounts for the effects of both low-level adaptation mechanisms and higher-level cognitive color processing. This iCAM06-based algorithm provides a principled approach to image splitting for HDR displays.

Figure 3 shows side-by-side comparison between images produced by the luminance square root and the iCAM06-based algorithms. Figures 3a and 3b show the front plane images created for the LCD display. These images represent tone-mapped SDR versions of the original HDR input. Note that the iCAM06 processed image has better contrast and is more colorful due to locally adaptive contrast compression and explicit handling of color and image appearance properties. Figures 3c and 3d show the
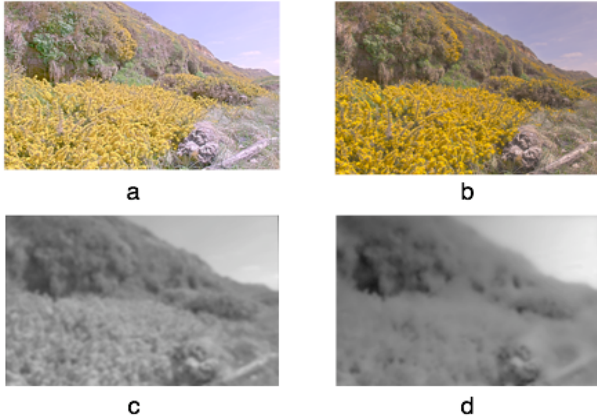
Figure 3: *(a) Square-root LCD image, (b) iCAM06-based LCD image, (c) Square-root backlight image and (d) iCAM06-based backlight image.*

achromatic images created for the DLP backlight. Although both images have the same number of pixels, the iCAM06-based image is much smoother and has greater perceived contrast. The smoothness is due to the fact that the local contrast adjustment in the iCAM06 model aims to preserve details for the LCD image. This also provides greater perceived contrast.

Experiments

To evaluate image quality and user preference for images processed using the square root and iCAM06-based HDR image splitting algorithms, we conducted a series of experiments that are described in the following sections.
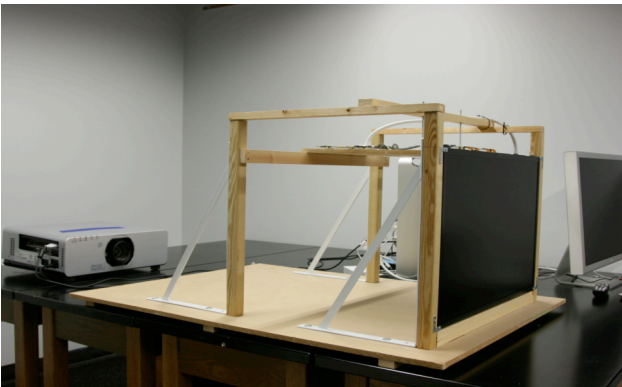

Figure 4: *HDR display used in the experiments. Note that the cover has been removed to show the components and normally stray light and interreflections are controlled.*

HDR display

To conduct studies of image quality on HDR displays, one needs an HDR display. While commercial HDR displays are slowly becoming available, they typically include "black box" image processing that does not allow for controlled testing of HDR image splitting algorithms. Therefore we constructed our own HDR display following the projector/LCD design described in [9]. The display (shown in Figure 4) is itself a modification of a system described in [19]. The display consists of a 30" Apple Cinema HD

LCD panel with the CCFL backlight removed, that is then backlit by a Panasonic PT-DW 6300 DLP projector. A Fresnel lens behind the LCD collimates the projector image, which is then passed though a diffusing sheet before entering the LCD. The LCD front plane has 2560x1600 addressable pixels and the DLP backlight has 1280x800 addressable pixels. Custom software [16] registers the two image planes geometrically and performs a full colorimetric characterization based on a modification of the Day et al. model [19] shown in Equation 1.

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \left(\begin{bmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{bmatrix} \times \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \begin{bmatrix} X_k \\ Y_k \\ Z_k \end{bmatrix} \right) \cdot (A(1-A_k)+A_k) \quad (1)$$

Here the 3x3 matrix represents the XYZ tristimulus values of the RGB primaries of the LCD. RGB and A represent radiometric scalars in the range 0 to 1 used to drive the LCD and DLP respectively, and the XYZ_k and A_k term account for the black level offset of the LCD and the minimum light output from the DLP projector.

The maximum luminance of the display is approximately 700 cd/m^2 with minimum luminance of 0.028 cd/m^2 for a small black region surrounded by a full white field. Under these near-worst-case conditions, the display dynamic range is 25,000:1. Higher values should be expected for less extreme conditions. The system is driven using custom MATLAB code that incorporates display control functions from the Psychophysics Toolbox [21].

Test images

An ideal HDR image splitting algorithm should be able to handle a wide range of source images. Therefore in testing, it is important to include HDR images with varying dynamic ranges and color gamuts and different categories of subject matter. To meet this requirement, forty images covering four content categories (Indoor, Day, Night, People) were selected from three HDR image databases: Fairchild's HDR image survey [22], standard HDR test images from Reinhard et al.'s High Dynamic Range Imaging book [23], and a selection of single and group portraits provided by Ward [24]. Thumbnails of the image set are shown in Figure 5.

Procedure

The experiments employed two psychophysical procedures: rating and pair comparison. In the rating study, observers were asked to rate the quality of a single image presented in the center of the screen on a scale from 1 to 5, where 1 was labeled "poor" and 5 was labeled "excellent". Overall observers rated 80 images (each image in the test set processed with both the square root and iCAM06-based algorithms). The images were presented in random order. To familiarize subjects with the experiment procedure, 10 practice trials were given with images randomly selected from the test set.

Pair comparison experiments also were conducted to directly assess user's preference for images processed by the two algorithms. The square root and iCAM06-based processed versions of an image were presented side by side on the screen and observers were asked to select the one they preferred with respect to different criteria. In five separate studies, observers were asked: a) Which image do you like better? b) Which image has better

Figure 5: *Thumbnails of the HDR images used in the experiments.*

colors? c) Which image looks sharper? d) Which image has better contrast? and e) Which image looks more natural? The image presentation sequence and their left/right positions on the screen were completely randomized.

Observers viewed the display in a dark room from approximately 60 cm. At this distance, each 780 pixel-wide image subtended approximately 12 degrees of visual angle horizontally.

Twenty-three observers participated in the studies. All were university staff or students. Both genders were represented and ages ranged from 23 to 66. The observers had varying cultural backgrounds and varying experience with imaging technology. All had normal or corrected to normal acuity and self reported normal color vision.

Results and discussion

Rating experiment

The average quality ratings given to the images processed by the square root and iCAM06-based algorithms are summarized in Figure 6. Note the similarity of the scores for the two algorithms and for the different images in the set. This similarity is further illustrated in Figure 7, which plots per image differences in the average ratings given for the two algorithms. An ANOVA was conducted to determine if the ratings for the two algorithms were significantly different, failed with a p-value of 0.7466. Thus it appears that overall, observers feel that the quality of the images produced by the two algorithms is comparable.

However, it may be that the variety of the images included in the overall result is masking trends within the data. To investigate this possibility, we plotted the mean ratings given to the images for each algorithm against measures of image dynamic range, average
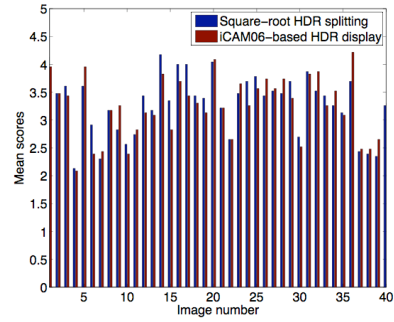

Figure 6: *Average quality ratings given to each image processed by the square root and iCAM06-based algorithms. Note the similarity of the per image ratings.*
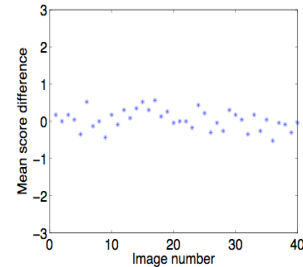

Figure 7: *Per image differences between the average ratings given for the two algorithms.*
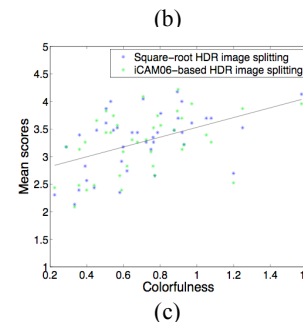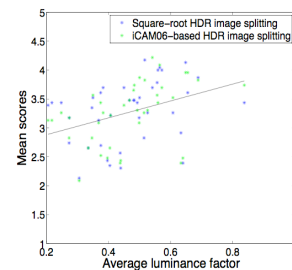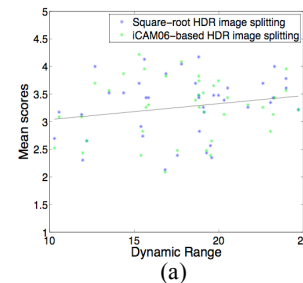

(a)

(b)

(c)

Figure 8: *Regression fits to mean image ratings as a function of a) image dynamic range; b) average image luminance factor; c) image colorfulness.*

luminance factor, and colorfulness. Equations (2-5) used to calculate dynamic range and average luminance factor are taken from [25]. Here L_w is the world luminance for each pixel, N is the total number of pixels in the image, and δ is a small value to avoid singularity. The colorfulness calculations are taken from the image appearance component of the iCAM06 model.

$$DynamicRange = \log_2 L_{Max} - \log_2 L_{Min} \qquad (2)$$

$$\bar{L}_w = \exp\left(\frac{1}{N}\sum_{x,y}\log(\delta + L_w(x,y))\right) \qquad (3)$$

$$AverageLu\min ance = \log_2 \bar{L}_w - \log_2 L_{Min} \qquad (4)$$

$$AverageLum.Factor = \frac{AverageLu\min ance}{DynamicRange} \qquad (5)$$

The resulting graphs are shown in Figure 8 with regression lines fit to the data. Although the lines show positive slopes, the low values of the coefficients of determination R^2 (shown in Table I) suggest that there are no meaningful trends in these data.

Table I: parameters of the regression fits

Image property	Slope	Intercept	R^2
Dynamic range	0.03	2.75	0.05
Avg. lum. factor	1.46	2.59	0.15
Colorfulness	0.09	2.64	0.24

We then investigated whether the kind of image processed by an algorithm was related to the ratings given by the observers and if there were any differences between the algorithms along this dimension. Figure 9 shows the mean ratings given for four image categories (Day, Night, Indoor, and People) for the two algorithms. An ANOVA (shown in Table II) was conducted to test if there were significant differences between the ratings given for the two algorithms given for the different categories. None of the tests showed significant trends, however it should be noted that the numbers of samples in the different categories was small, so the reliability of the statistical tests is questionable. We are current collecting more data to further investigate the issue of categorical differences between the algorithms with respect to quality ratings.
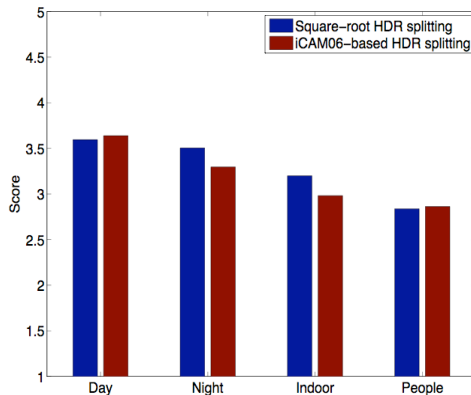

Figure 9: *Mean score of images for different categories.*

Table II: test statistics for differences in image ratings

Categories	Day	Night	Indoor	People
p-value	0.80	0.09	0.36	0.87

Paired-comparison experiments

Unlike in rating experiments, in paired-comparison experiments, observers are forced to compare one stimulus to another. This approach can reveal differences in sensitivity or preference that may be masked by the rating method. The results of five paired comparison preference experiments are summarized in Figure 10. Each pair of bars corresponds to the results for each of the preference criteria tested (overall preference, better color, more natural, better contrast, and better sharpness). The heights of the bars indicate how frequently the square root and iCAM06 processed images were preferred.
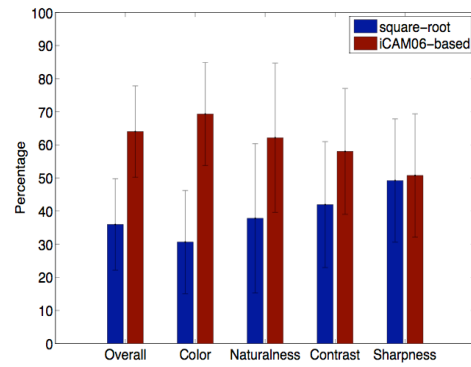

Figure 10: *Paired comparison results on all five attributes.*

Note that for all five criteria, the iCAM06-based processed images were preferred over the square root images. However there was significant individual variation in preference, so while the preference trends for naturalness, contrast, and sharpness favored the iCAM06-based algorithm, the differences were not statistically significant. On the other hand, the iCAM06-based algorithm was significantly preferred overall and with respect to color. We believe this can be attributed to the improved color processing incorporated in the iCAM06-based algorithm.

To complement the image category analysis we performed on the rating data, we also split the preference results by image category. These data are summarized in Figure 11.

The graphs show some interesting trends. First, different kinds of images show various amounts of preference difference, with the Day and People image sets showing strong preference for the iCAM06-based algorithm, and the Night and Indoor groups showing similar trends but no statistically significant differences. Second, the Night and Indoor images actually showed slight preference for the square root algorithm with respect to contrast and sharpness. Again, the statistical significance of these results should be evaluated with the understanding that the image sample sizes in the categories were relatively small. More extensive studies are ongoing.

To get deeper insight into the effects of image content on observer preference for images processed by the two algorithms, we ranked the preference scores and identified the least and most preferred images for the two methods. The most preferred images
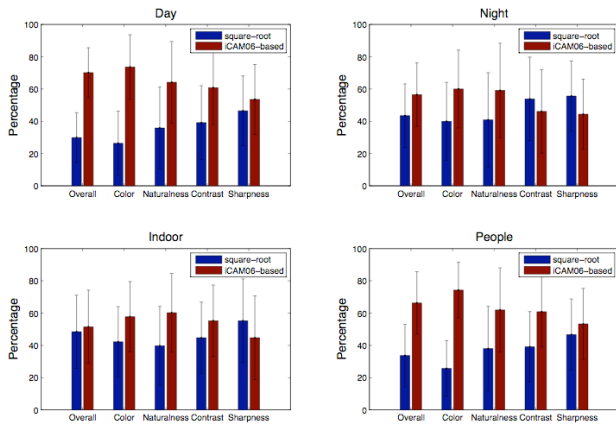
Figure 11: Paired comparison results for different image categories.

in each set are shown in Figure 12. The top row shows the most preferred images for the square root algorithm (these are consequently the least preferred for the iCAM06-based algorithm). Note that these images are generally dark, have high dynamic range from concentrated light sources and have highly saturated colors. On the other hand, the most preferred images for the iCAM06-based method are shown in the bottom row. These are generally bright, daylit, and show natural environments. We believe that the enhanced color processing provided by the iCAM06-based algorithm may explain this pattern of preference differences, with color saturation playing a role, but further investigations are necessary before any conclusions can be drawn.
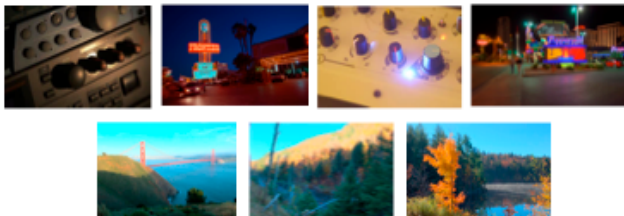

Figure 12: Top row: most preferred square root processed images; Bottom row: most preferred iCAM06-based processed images.

Conclusion

In this paper we have introduced a new appearance-based algorithm for HDR image splitting. The algorithm incorporates the iCAM06 image appearance model and seeks to create displayed HDR images that faithfully represent the appearances of HDR scenes. To evaluate the algorithm, we ran a series of experiments that compared images processed by the new algorithm with images processed by the widely used luminance square root algorithm. Although a single image quality rating study showed little difference in quality ratings, pair comparison preference studies showed that observers preferred images produced by the iCAM06-based algorithm overall and with respect to color rendering. Both studies showed differences in observer ratings and preferences for Day, Night, Indoor, and People image categories. This work is preliminary and much remains to be done, but the new appearance-based algorithm introduced in this paper represents a promising and principled approach to image splitting for HDR displays.

References

[1] J.A. Ferwerda, S. Pattanaik, P. Shirley, and D.P. Greenberg, A model of visual adaptation for realistic image synthesis. ACM Transactions on Graphics (Proceedings SIGGRAPH '96), pg. 249-258, (1996).

[2] J. Kuang, High-dynamic-range digital photography: image display and perception. Unpublished Ph.D Thesis. Rochester Institute of Technology, (2006).

[3] S. Mann, and R.W. Picard, Being undigital with digital cameras: extending dynamic range by combining differently exposed pictures, Proceedings IS&T, pg. 442-448 (1995).

[4] P. Debevec and J. Malik, Recovering high dynamic range radiance maps from photographs, Proceedings ACM SIGGRAPH '97, pg. 369-378. (1997).

[5] M. Robertson, S. Borman and R. Stevenson, Dynamic range improvements through multiple exposures. Proceedings of the International Conference on Image Processing (ICIP) '99, pg. 159-163. (1999).

[6] T. Mitsunaga and S. K. Nayar, Radiometric self calibration. Proc. IEEE CVPR '99, pg. 472-479. (1999).

[7] P.M. Acosta-Serafini, I. Masaki, C.G. Sodini, Single-chip imager system with programmable dynamic range, U.S. patent 6,977,685, (2005).

[8] S. B. Kang, M. Uyttendaele, S. Winder and R. Szeliski, High dynamic range video, ACM Transactions on Graphics (Proc. SIGGRAPH '03), pg. 319-325. (2003).

[9] Helge Seetzen, Wolfgang Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, A. Vorozcovs, High dynamic range display systems. ACM Transactions on Graphics (Proc. SIGGRAPH '04) pg. 760-768. (2004).

[10] H. M. Visser, J. J. W. M. Rosink, N. Raman, and R. Rajae-Joordens, Tuning LCD displays to medical applications, Proc. EuroDisplay '05, pg. 74–77, (2005).

[11] P. Ledda, G. Ward, and A. Chalmers, A wide field, high dynamic range, stereographic viewer, Proc. ACM GRAPHITE 2003, pg. 237–244, (2003).

[12] O. Bimber, Weimar and D. Iwai, Superimposing dynamic range. ACM Transactions on Graphics (Proc. SIGGRAPH '08), pg. 1-8. (2008).

[13] J. Kuang, G.M. Johnson, M.D. Fairchild. iCAM06: A refined image appearance model for HDR image rendering. J. Visual Communication and Image Representation. 18(5), pg. 406-414. (2007).

[14] Greg Ward, Maryann Simmons. JPEG-HDR: a backwards-compatible, high dynamic range extension to JPEG. 13th Color imaging conference.

[15] M. Trentacoste, W. Heidrich, L. Whitehead, H. Seetzen and G. Ward, Photometric image processing for high dynamic range displays. J. Visual Communication and Image Representation. 18(5) pg. 439-451. (2007).

[16] S. Luka and J.A. Ferwerda, Colorimetric image splitting for high dynamic range displays. Proceedings SID 2009 Annual Conference, pg. 1298-1301. (2009).

[17] G. Guarnieri, L. Albani, and G. Ramponi, Minimum-error splitting algorithm for a dual layer LCD display—Part I: Background and theory J. Display Technol., 4(4), pg. 383-389, (2008).

[18] G. Guarnieri, L. Albani, and G. Ramponi, Minimum-error splitting algorithm for a dual layer LCD display—Part II: Implementation and results, J. Display Technol., 4(4), pg. 391-397, (2008).

[19] J.A. Ferwerda and S. Luka, A high resolution high dynamic range display for vision research (abstract), 8th Annual Meeting, Vision Sciences Society, Journal of Vision, 9(8), 346a, (2009)

[20] E.A. Day, L. Taplin, R. S. Berns, Colorimetric characterization of a computer-controlled liquid crystal display, Color Res Appl, 29, pg. 365-373, (2004).

[21] D. H. Brainard, The psychophysics toolbox, Spatial Vision 10, pg. 433-436, (1997)

[22] M.D. Fairchild. HDR photographic survey, http://www.cis.rit.edu/fairchild/HDR.html, (2008).

[23] E. Reinhard, G. Ward, S. Pattanaik and P. Debevec, *High dynamic range imaging: acquisition, display, and image-based lighting*. Elsevier/Morgan Kaufmann, Amsterdam, (2006)

[24] Personal communication with Greg Ward. (2010)

[25] J. Kuang, G.M. Johnson, M.D. Fairchild. Testing HDR image rendering algorithms. Proc. IS&T/SID 12th Color Imaging Conference. (2004)

Author Biography

Dan Zhang is a graduate student in the Color Science Master program at Rochester Institute of Technology. She received B.S. degree in Electronic Science and Technology from Yanshan University and M.S. degree in Optical Engineering from Beijing Institute of Technology.