Visualization of High Dynamic Range Images

Alvaro Pardo* Guillermo Sapiro†

Abstract

A novel paradigm for information visualization in high dynamic range images is presented in this paper. These images, real or synthetic, have luminance with typical ranges many orders of magnitude higher than that of standard output devices, thereby requiring some processing for their visualization. In contrast with existent approaches, which compute a single image with reduced range, close in a given sense to the original data, we propose to look for a representative set of images. The goal is then to produce a minimal set of images capturing the information all over the high dynamic range data, while at the same time preserving a natural appearance for each one of the images in the set. A specific algorithm that achieves this goal is presented and tested on natural and synthetic data.

EDICS: 4-OTHD.

1 Introduction

High dynamic range images contain a wide range in luminance, many times in the order of tens of thousands different values.[1] These images could be natural, obtained for instance from multi-exposure photographs [1] or with a multiple exposure sensor [6, 14], or synthetic, in the case of computer graphics applications. These images have ranges that greatly exceed that of the output device. The question is then how can we *reproduce* and *visualize* such images in a standard output device. We address this problem in the present paper.

*IMERL & IIE, Facultad de Ingeniería, Universidad de la República, CC-30, Montevideo, Uruguay. apardo@iie.edu.uy

†Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455. guille@ece.umn.edu

[1] These images are also referred in the literature as *high contrasted*.

Before proceeding, let us introduce some basic terminology. The *scene* is the real or synthetic picture we perceive without involving any output device between it and our eyes.[2] An *image*, on the other hand, is what we can see using the output device, or its internal computer representation as an array of digital values. The key problem is how to translate from scenes to images while preserving the relevant scene information, producing a natural looking image, and avoiding common artifacts such as *halos*, which are due to local gradient reversals [2].

Numerous applications exist for high dynamic range images. One is computer graphics and the production of synthetic images with realistic or hyper-realistic appearance. Another application covers high dynamic range photographs, which are able to capture much more detailed scene information than standard photographs. Recently, methods to acquire such photographs have been developed [1, 14, 6]. In particular, in [14, 6] the described system is already included in a camera prototype. These systems make it possible to capture a highly detailed range representation of the scene and later process the data in order to select the image/s that better fulfills the given requirements. These images could also improve computer vision and image analysis algorithms that usually rely on limited range data. This is particularly relevant in scenarios where we do not have complete control over the illumination, like medical applications for instance. Therefore, there is a need to develop algorithms to perform the translation from scenes to images, algorithms as the ones discussed and presented in this paper.

When addressing this translation problem we have to consider the digital nature of output devices. Indeed, this is a very important issue when trying to display details. Somehow, the problem of high dynamic range is a resolution problem; with a small number of output levels we have to display a highly detailed image. Things become even more problematic when the number of output levels is less than the minimum needed to obtain a good *quantization*.

We can classify the existing approaches for the translation from scene to image in two main groups (examples of these will be detailed in the next section). The first group consists of algorithms that map the original range to the output range while attempting to preserve the subjective perception of the scene, e.g., [5]. Although this idea of "tone mapping" works quite well, it has

[2] In the case of synthetic scenes, it might be sometimes very difficult to directly visualize the data.

some caveats. First, it is not able to reproduce all the details present in the scene. Second, the method breaks down when the input range is too wide compared with the available output range.

In the second group, we have algorithms that favor the visualization of details instead of the subjective perception of the scene. As examples of this approach we have the works of Tumblin and Turk [13] and DiCarlo and Wandell [2]. Both apply a multiscale decomposition to discriminate between illumination and details. The main problem with this idea is that, although correct in theory, it usually introduces halos in the output image. The local mappings produced by these techniques violate the basic monotonicity principle, the pixel value order is not necessarily preserved and darker (brighter) regions in the scene might become brighter (darker) in the image. Moreover, these approaches tend to have a large number of parameters, generally hard to control in an automatic fashion.

Our proposed paradigm attempts to have the best of both groups mentioned above. We propose a method which captures the image details while preserving the natural appearance of the scene. As we will explain below, there is an intrinsic limitation in representing a high dynamic range image with only *one standard* output image. Sometimes it is practically impossible to find an output image containing all the relevant information in the high dynamic range image that represents the scene. For this reason, we argue for a method to obtain a *set of images* containing all the relevant information of the original scene and displayed in a suitable way.

1.1 Related work

In this section we review some of the recent literature on the reproduction of high dynamic range images. First, we address the methods that attempt to preserve the subjective perception of the scene and then the ones that favor the visualization of details.

In [12], Tumblin and Rushmeier developed a tone mapping operator using models of human perception. The main drawback of their algorithm is that they use a global brightness adaptation, dark and bright regions are clipped. Schlick [10] concentrated on a simple method for computing a local tone mapping.

Ferweda et. al. [4] noted the connection between light levels, color, and acuity. Using this work, Larson et. al. [5] proposed a global tone mapping operator which adjusts the histogram of the scene based on psychophysical models for color, glare, and acuity perception. The results of this simple and elegant method, which reduces to just a global map, are very good and with high fidelity to the subjective perception of the scene. The problem is, as for most of the methods in this category, when the input range is too large. Note that since this method tries to preserve the original perception of the scene, details that are hard to see in the original scene will be difficult to see in the output image as well.

In [11] the authors introduce two methods to display high dynamic range images. The first one is intended to display synthetically generated images. The image is decomposed into layers of lighting and surface properties. The light layer, which contains most of the high contrast, is compressed and added back to the surface layers containing details and texture. In this way, high contrast is reduced while preserving the details and texture from the original image. This method is designed to work only for synthetic images, for which the different layers are easily computed. (We will come back to this point later when reviewing the work of Tumblin and Turk [13].) For natural scenes, they proposed a locally adaptive method, denoted as the foveal display, which is inspired by eye movements. The user selects a point of attention and the algorithm computes an output image with preserved contrast in the foveal region (a region around the selected point). It is important to remark that this approach is dynamic in the sense that a set of images is generated with the aid of user interaction. The drawback with this approach is that the user needs to select a point of attention. Moreover, the user should be able to "see" everywhere in the image to choose the points of attention. Problems could then arise if the image presented to the user is not adequate for inspection. On the other hand, we should note that this approach is connected to ours in the sense that we also suggest to compute a set of images instead of a single one.

The work of Pattanaik et. al. [7] proposes a multiscale model for the representation of pattern, luminance, and color in the human visual system. The main problem with this work is that, although interesting for its detailed modeling of the human visual system, it cannot avoid halos.

Continuing with the idea of segregating the image into layers of lighting and details [11], Tumblin and Turk [13] proposed a multiscale approach to extract a hierarchy of details and boundaries. Their idea is to mimic the way artists work from coarse to fine when recreating highly contrasted scenes in low contrasted mediums. Artists start with a sketch of strong features and progressively add small details. They propose to decompose the image into strong and weak features using a multiscale operator, then, only strong features are compressed. Although the idea is very attractive, their algorithm cannot avoid halos completely and, as pointed out by the authors, needs the difficult tuning of some crucial parameters.

In our view of the problem, the two groups of works just described are not completely equivalent, they basically address slightly different problems. Furthermore, although similar, their solutions cannot be easily compared, since one representation tries to capture the subjective appearance of the scene under the limitation of the output/display device, while the second group attempts to preserve the scene details. While enhancing details, we may be adding information not perceptually present in the original scene. On the other hand, while preserving visual appearance, details might be omitted.

Let us conclude by pointing out that recently, the authors in [9] have shown how to obtain panoramic images with high dynamic range.

1.2 Our contribution

The approaches described above produce a single image per scene or per focused scene region. It is a quite optimistic assumption that we can accurately represent and visualize the information of an image with tens of thousands of values with just a few hundred. This is what motivates the paradigm here proposed, meaning the use of a set of images to represent such highly detailed information. We could say that while the algorithms described above deal with the *reproduction* of the scene, the technique here proposed deals with its *visualization*. Moreover, we argue that not only the set of images has to accurately visualize the relevant information present in the scene, but also has to do it in a visually pleasant form. We present a particular algorithm to exemplify this new paradigm. In the proposed algorithm, each image is produced by a different monotonic global map, thereby avoiding gradient reversals typical of the local schemes. The locality is achieved by letting this global map "stretch" different regions for each one of the images in the set.

2 The method

Our idea is to propose a simple and effective method to visualize *all* the information in the scene in a *pleasant* way. *All* means that we would like to capture as many details as possible, and *pleasant* means a procedure which appears natural to the observer. Finally, we address the problem of halos and other artifacts in the solution.

Let us assume we have a scene with dark and bright areas, and details all over it. If we want to visualize all the details we need first to have a minimum resolution available (number of output levels), and second to be able to "see" in every region. In order to understand what is meant by "see," we present a couple of simple examples. Consider the image of a dark room, in order to "see" the details we would usually turn on the lights. Now, suppose we are in the beach and everything is too bright, in this case to "see" we would probably wear sunglasses to reduce the amount of light. In both cases, the information is out there but it cannot be seen. In photographic terms, this is due to under-exposure or over-exposure.

These problems are not at all new, especially in photography they are of great importance. If we take a picture of an object and there is light coming from behind it, we will not capture all its details. To solve this, photographers put extra light to the front of the object. We are going to borrow this idea and modify the luminance of the scene to capture all the details over it. All the proposed operations are simply contrast changes. In this way, artifacts such as halos are not introduced.

The first observation is that if we illuminate a given region, we will be stretching its output range, thereby using more output levels. On one hand we display this region with good light and resolution, while on the other hand we might be compressing and missing details in other regions. Hence, there is clearly a competition between the output fidelity of different regions, and unfortunately, it is difficult or impossible to find a satisfactory solution with a single output. To overcome this, we propose to generate a sequence of images with different resolutions in each scene region (space varying resolution). This sequence could be either observed as a movie or as a set of still images could be extracted from it. The key point here is that for many applications more than one output image is a reasonable solution.

The basic idea is then to distribute the existing resources among different output images. In the case of only two regions, we first display dark regions adding extra light to them and then we slowly swap resources to bright areas. This simple idea resembles the control of illumination during acquisition and is perceived as natural by human observers.

2.1 Outline of the algorithm

Before presenting our proposed algorithm, let us give some basic notation:

(r, g, b)	Input color primaries
L	Input luminance
$[L_{wmin}, L_{wmax}]$	Input luminance range
L'	Modified luminance
$[L_{dmin}, L_{dmax}]$	Output luminance range
(R, G, B)	Digital output values

We are now ready describe the different steps of the algorithm.

1. **Compute image luminance:**

   From the (r, g, b) primaries compute the luminance L (in cd/m^2) and the color information $(r/L, g/L, b/L)$. We process the luminance while preserving the color information.

2. **Segment the image:** Divide the image into two or more regions of interest. This is achieved splitting the histogram into sub-intervals $[L_{wmin}, L_1], [L_2, L_3], ..., [L_n, L_{wmax}]$. If we segment the image arbitrary we could loose the monotonic restriction of the tone mapping. We then choose a simple procedure that segments the image into bright and dark areas to illustrate the idea, see section 3.1. From now on then, in this section, we assume two regions, $[L_{wmin}, L_{w^*}]$ and $[L_{w^*}, L_{wmax}]$.

3. **Modify the luminance:** Apply Larson's et al. histogram adjustment algorithm, [5] (see Appendix), to each interval. Map $[L_{wmin}, L_{w^*}]$ to $[L_{dmin}, L_{d^*}]$ and $[L_{w^*}, L_{wmax}]$ to $[L_{d^*}, L_{dmax}]$. The important point here is that when using the human contrast sensitivity function, the

mapping does not produce a contrast greater than the one present in the original scene. Since we are working in regions, this means controlling the contrast over that given region. This step is not a traditional histogram adjustment since it modifies the output range and the distribution within it.

It is clear that if we select L_{d^*} close to L_{dmax} we will be visualizing the dark areas, $[L_{dmin}, L_{d^*}]$, with a wider range than the bright ones. Thus, by changing L_d^* we modify the resources assigned to each interval of the original luminance. A first solution to the problem of visualization high dynamic range data is to create a movie by increasing L_{d^*} from L_{dmin} to L_{dmax}. This gives us a sequence, which starting from the image with all resources allocated to the dark areas, slowly moves to an image with all resources allocated to the bright areas. This is just a nice way of visualizing all the information. A second possibility is to choose just a certain fix number of images. We discuss below how to select these images.

In the case of three intervals or more, the idea is the same. We start with most resources allocated to the first interval and we swap them to the next interval to the right. That is, we first split the output luminance range into three intervals, $[L_{dmin}, L_{d1}], [L_{d1}, L_{d2}], [L_{d2}, L_{dmax}]$, with $\Delta_M = L_{d1} - L_{dmin}$ and $\Delta_m = L_{d2} - L_{d1} = L_{dmax} - L_{d2}$. Initially, Δ_M is much greater than Δ_m. Then we swap resources from the first interval to the second one, obtaining the intervals $[L_{dmin}, L_{d1} - \delta], [L_{d1} - \delta, L_{d2}], [L_{d2}, L_{dmax}]$. We iterate this procedure until the step k when $L_{d2} - L_{d1} - k\delta = \Delta_M$, and then we proceed to do the same with the second and third intervals.

4. **Quantization:** Quantize and gamma correct the recomputed primaries $(\frac{r}{L} * L', \frac{g}{L} * L', \frac{b}{L} * L')$ to obtain the digital output values (R, G, B).

We should note that before displaying the images we correct the color in the mesopic and scotopic range, where our perception loses efficiency, using the radiance library [8].

2.2 Information assessment

As we mentioned before, the evaluation of the output image is mostly subjective. However, it would be of interest to have an automatic procedure to assess the information content of each image in the set. With it, we will be able to extract the best or set of best images in the sequence.

If we consider each image as a message, its information can be measured with the entropy function. We can therefore use this as a plausible measure of the image information. From the histogram of the output image we find the probabilities of each output level, p_i, and with them we compute the entropy as

$$H = -\sum_i p_i \log p_i$$

Note the relation between histogram equalization and entropy. The entropy is maximum when all symbols are equally probable, which implies a flat histogram. Additionally, entropy maximization captures our subjective preference towards well-contrasted images.

3 Estimation of the number of images

The minimal number of images needed for a satisfactory visualization of all the details in the high dynamic range data depends of course on the particular image being processed. Indeed, the number of output levels depends on the luminance distribution within the given image. Therefore, to obtain an estimation of the number of images we use the number of clusters in the luminance histogram.[3] To display the image without loss of information we would need, roughly, as many output levels as clusters present in the image. However, not all output levels can be used if we want to distinguish between them. That is, we should use less output levels in order to guarantee that different input levels are mapped into visually different output levels. Methods based on the Weber's law of contrast perception are unfortunately too pessimistic (in our experience, for a greyscale image only one fourth of the output levels could be used). Hence, we use an empirical rule of thumb that says that we can allocate around 200 clusters per (8-bits) image in the set. This is just a coarse estimation that works fairly well in practice.

[3]We compute the number of clusters using the algorithm for peak detection proposed in [3].

Having an estimation of the number of output images required to satisfactory visualize the scene information, a segmentation is needed in order to find the location of the luminance intervals. Obtaining a good segmentation in this case is a difficult problem. To begin with, we want to visualize information even in places where there is not enough (or too much) light, and standard segmentation algorithms will fail to segment these regions properly since bright and dark spots will be typically assigned to a single region. Second, it is hard to obtain segmentations with regions matching our perception. We then select a simple approach based on the histogram. If needed, the user can correct the segmentation according with his/her needs. In what follows, we explain this segmentation procedure.

3.1 Histogram-based segmentation

The segmentation method we propose to use has been satisfactory used for documents in order to separate text from background. The text is made of dark pixels while the bright ones represent the background. To segregate the text from the background, a threshold t that maximizes the sum of the entropy of both distributions is computed. Given the probabilities p_B of dark pixels and p_W of bright pixels, t is computed maximizing the entropy

$$H = -\sum_{i=0}^{t} p_B \log p_B - \sum_{i=t}^{255} p_W \log p_W$$

In our case we want to separate dark and bright regions. Using the histogram of the luminance and computing the above probabilities we find the threshold t that optimally segments dark and bright areas. If we use only two images, we only need to apply this procedure once, otherwise,we iterate this scheme.

Using the estimation of output images and the segmentation algorithm described above we split the image into regions. This procedure can be seen as a binary tree construction, where in each step an interval is divided in two. At the end, the user can select the important leaves from it if user interaction is constructed in the process.

4 Results

We now present a number of examples of our proposed algorithm. For the examples here presented, no user intervention was needed and the number of images and segmentation was performed as explained above. Movies showing the full set for the examples below can be found at www.ece.umn.edu/users/guille/hdr.html.

The first example is shown in Figure 1. This high dynamic range image was captured with Devebec's method [1]. Before discussing the results obtained for this example, we give specifics on the parameters of the algorithm.

The first step is the segmentation of the original image (see Figure 1). We found 419 clusters in the image luminance and therefore we use three output images (since we consider no more than 200 clusters per image). The intervals where found applying twice the algorithm described in 3.1. First, we obtained a segmentation of bright and dark areas and then we sub-segmented again the dark interval. Details of the segmentation are given in Table 1. Second, we assign the given resources to each region. In this case we assumed $L_{dmax} = 100\ cd/m^2$ and $L_{dmin} = 2\ cd/m^2$, and selected initial ranges $\Delta_M = 86\ cd/m^2$ and $\Delta_m = 6\ cd/m^2$. For the next image in the set, the resources are re-allocated moving them from left to right in steps of $\delta = 5\ cd/m^2$ (see the above mentioned web page for the full set of images obtained).

Region	Range(cd/m^2)	% Percentage of the total range
0	[0.12,2.27]	5.32e-3
1	[2.27,83.09]	0.20
2	[83.09,40416.77]	99.79

Table 1: Information for Memorial church image.

From the data in Table 1 we can draw several conclusions. To begin with, the first regions have very small luminance values. This means that is very difficult, even being in front of the real scene, to perceive details there. The dynamic range for this region is extremely small, representing only the 5.32×10^{-3} % of the total input range. In other words, although region 0 covers a significant

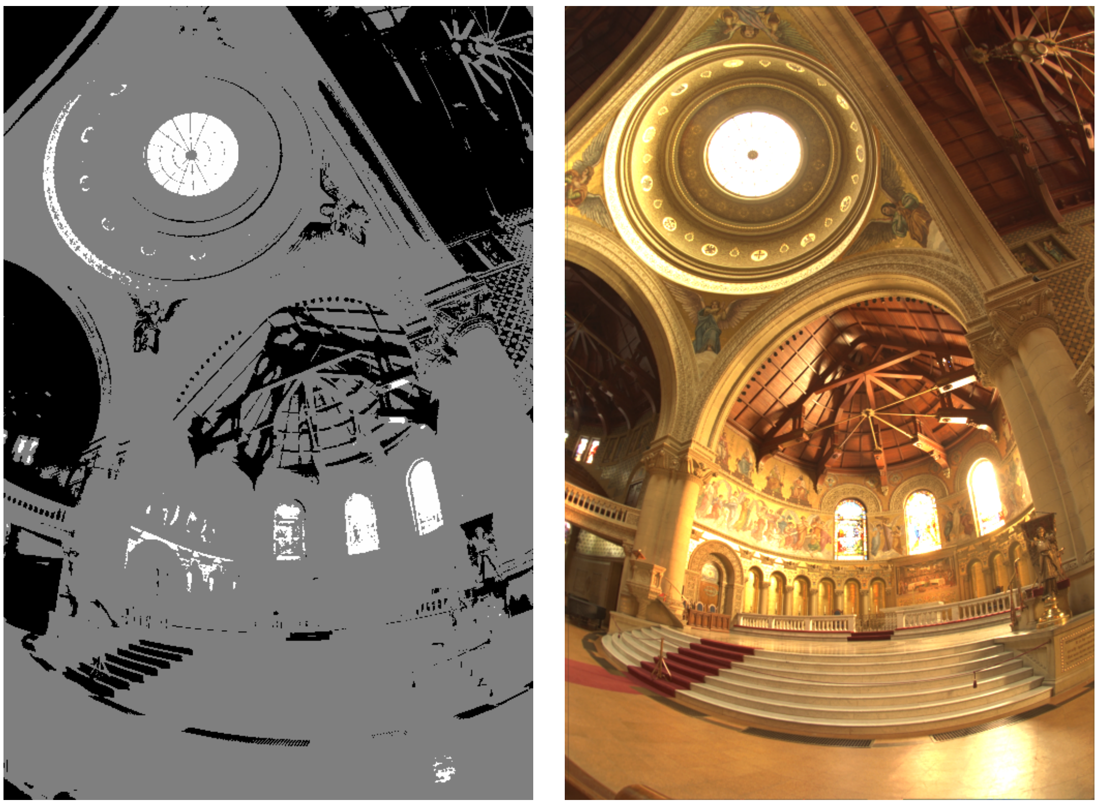

(a) (b)

Figure 1: (a) Segmented image into three regions. (b) The image with a linear tone mapping. Note how the segmentation based on the histogram provides a good partition, the dark and bright areas are clearly segmented. Also, note how a linear map saturates the bright areas.

part of the scene, it expands an insignificant range from the total. Once again, it is clear that high dynamic range data it is not just a problem of output range, it is in fact a problem of resolution within regions in the scene.

In Figure 2 we show the set of images obtained with our algorithm. Each image represents the image with maximal entropy per region. With these images we can see more details than in the single image processed with Larson's algorithm (Figure 3). The resulting images in Figure 2 are fairly natural, specially for the last two images. The three images selected visualize more information, especially in very dark and very bright regions. Hence, we managed to display the original scene in a set of images, making visible some information that was obscured in the original scene. Note for example the details in the dark corners and in the bright windows. In the first image observe the upper left and upper right corners, the details in the ceiling are clearly visible now. In the second image observe the walls. Finally, the third image represents the windows. Note how we make visible the details in the windows. These areas are completely lost with a linear tone mapping (simple linear map of the input range onto the output range), see Figure 1.

Since our method extends Larson's one, among all the images in the set, there should be one image close to the one obtained with Larson's algorithm. In other words, if we resign to use more than one image and visualize all the details in the scene, we should be able to extract a single satisfactory one from the computed set. To do that we take the image that maximizes the global entropy. If we compare it with the image computed with Larson's algorithm, see Figure 3, we can see that it contains enough details everywhere. This is due to the modification of the luminance distribution, we expanded regions that were originally very compressed. However, it is clear from the zoomed images, Figure 4 and Figure 5, that it does not capture dark and bright details as well as the images with maximum resources per region.

The same methodology was applied to the bathroom image. Here we selected two regions ($[3, 262]$ cd/m^2,$[299, 177222]$ cd/m^2). Figure 6 shows the segmented image together with the image obtained with a linear tone mapping. In this case, although the number of computed clusters indicated the use of three images in the set, we selected only two regions of the decomposition since they correctly represent the image. For this image the algorithm parameters are $\Delta_M = 91$ cd/m^2,

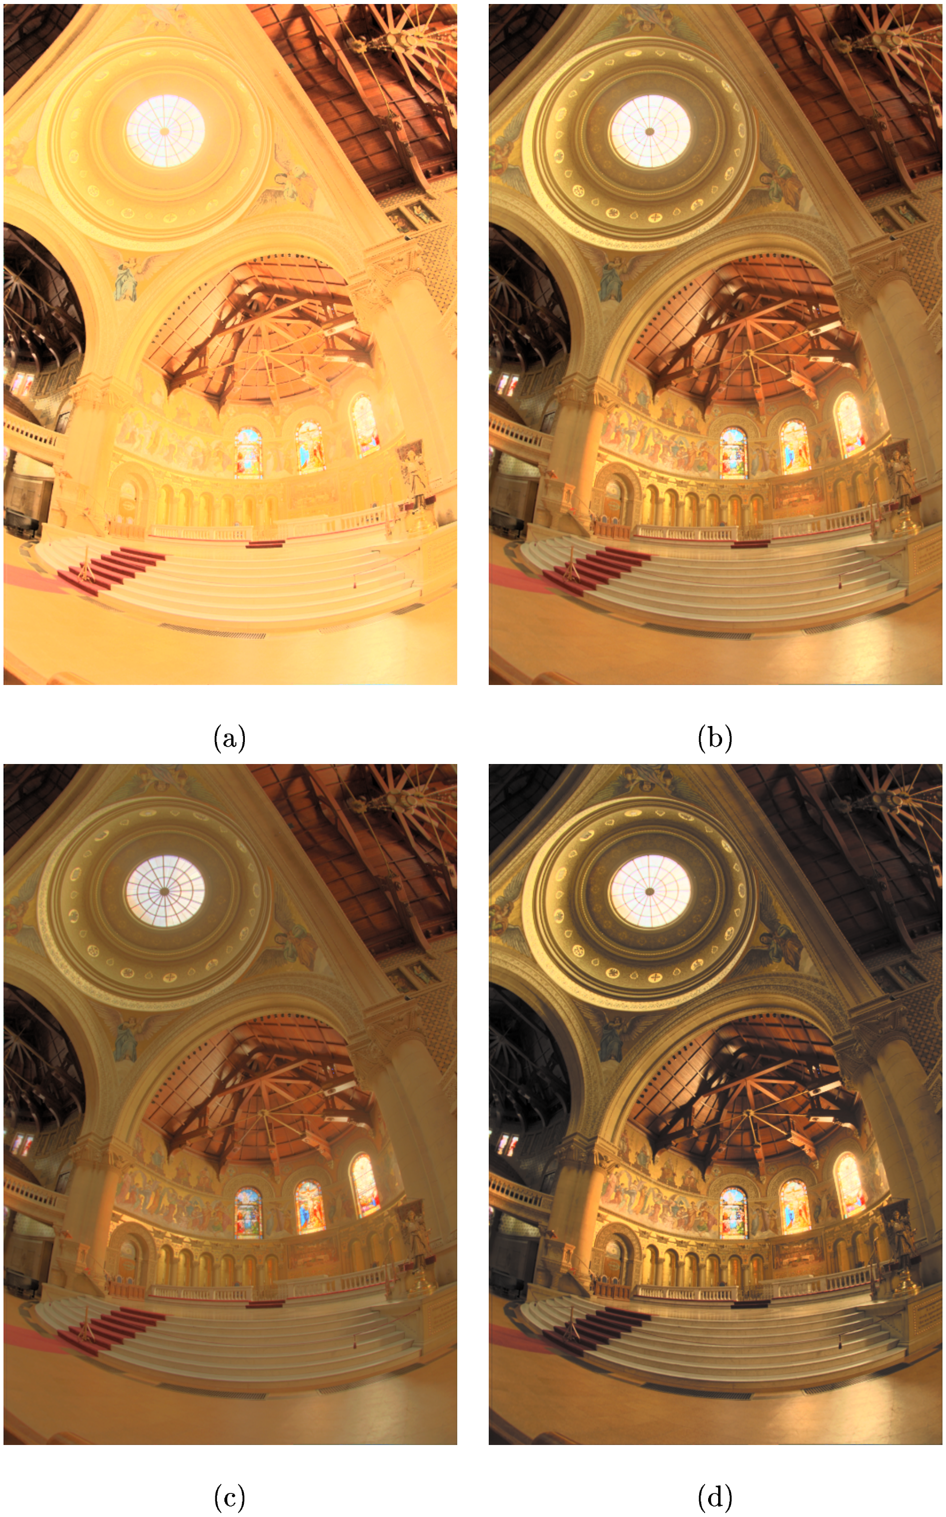

Figure 2: For each region we display the image with maximum entropy in the region. Images that visualize (a) region 0, (b) region 1 and (c) region 2. (d) Image obtained with Larson's algorithm.

(a) (b)

Figure 3: (a) Image processed with Larson's algorithm. (b) Image with maximum global entropy among all the images computed. Note how this image shows more details.

(a)

(b)

(c)

(d)

Figure 4: Zoomed images for: (a) Figure 2-(a), (b) Figure 2-(b), (c) Figure 2-(c), and (d) Figure 2-(d).

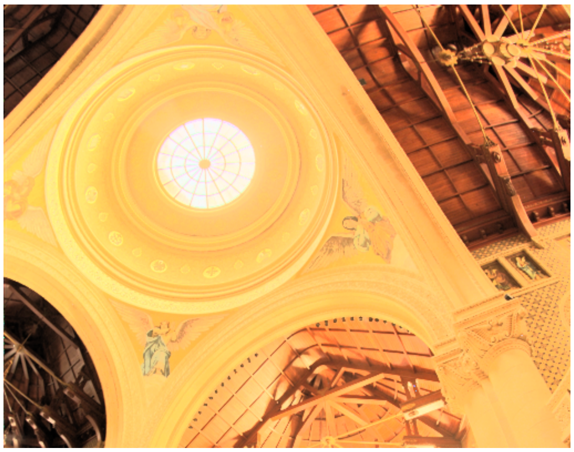

(a)

(b)

(c)

(d)

Figure 5: Zoomed images for: (a) Figure 2-(a), (b) Figure 2-(b), (c) Figure 2-(c), and (d) Figure 2-(d) computed with Larson's algorithm.

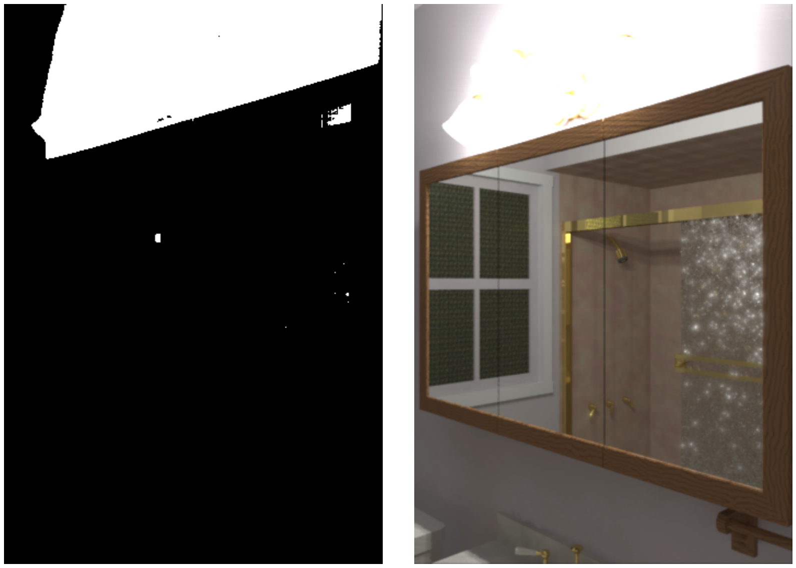

(a) (b)

Figure 6: (a) Segmented image. (b) Image with a linear tone mapping.

$\Delta_m = 7\ cd/m^2$ and $\delta = 4\ cd/m^2$. The image with all resources allocated to the bright region captures in a better way the details in the lamps while the one with all resources in the dark region accurately captures details in the rest of the scene, see Figure 7 and Figure 8. Finally, the image with maximum global entropy balances both regions and is comparable with the image obtained by Larson's algorithm, Figure 7.

The last example is the office image. The segmentation with intervals $[0.22, 7.43]\ cd/m^2$, $[7.43, 406.50]\ cd/m^2$ is shown in Figure 9 together with a linear tone mapping. The details in the floor and in the wooden desk are more clear in the top image on Figure 10, this details are difficult to visualize in the bottom image. Finally, once again, the image with maximum global entropy offers a balance between both images and outperforms in this respect the image obtained with Larson's algorithm.

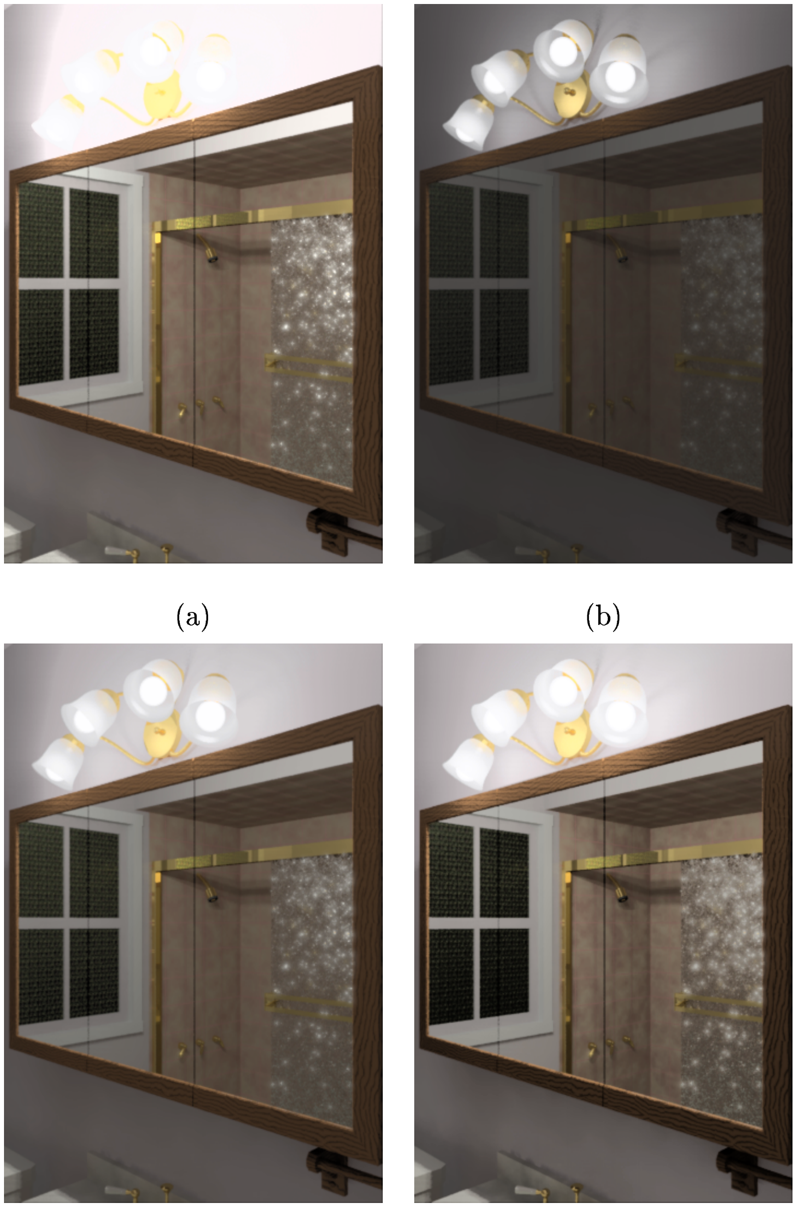

(a) (b)

(c) (d)

Figure 7: For each region we display the image with maximum entropy in the region. Images that visualize: (a) region 0 and, (b) region 1. (c) Image with maximum global entropy. (d) Image obtained with Larson's algorithm

(a)

(b)

(c)

Figure 8: Zoomed images for: (a) Figure 7-(a), (b) Figure 7-(b), and (d) Figure 7-(d).

(a)

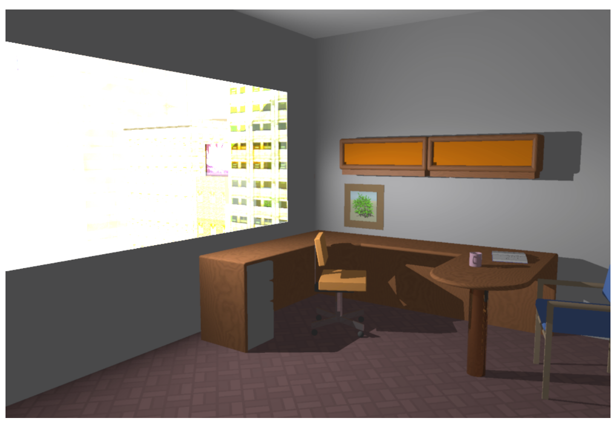

(b)

Figure 9: (a) Segmented image. (b) Image with a linear tone mapping.

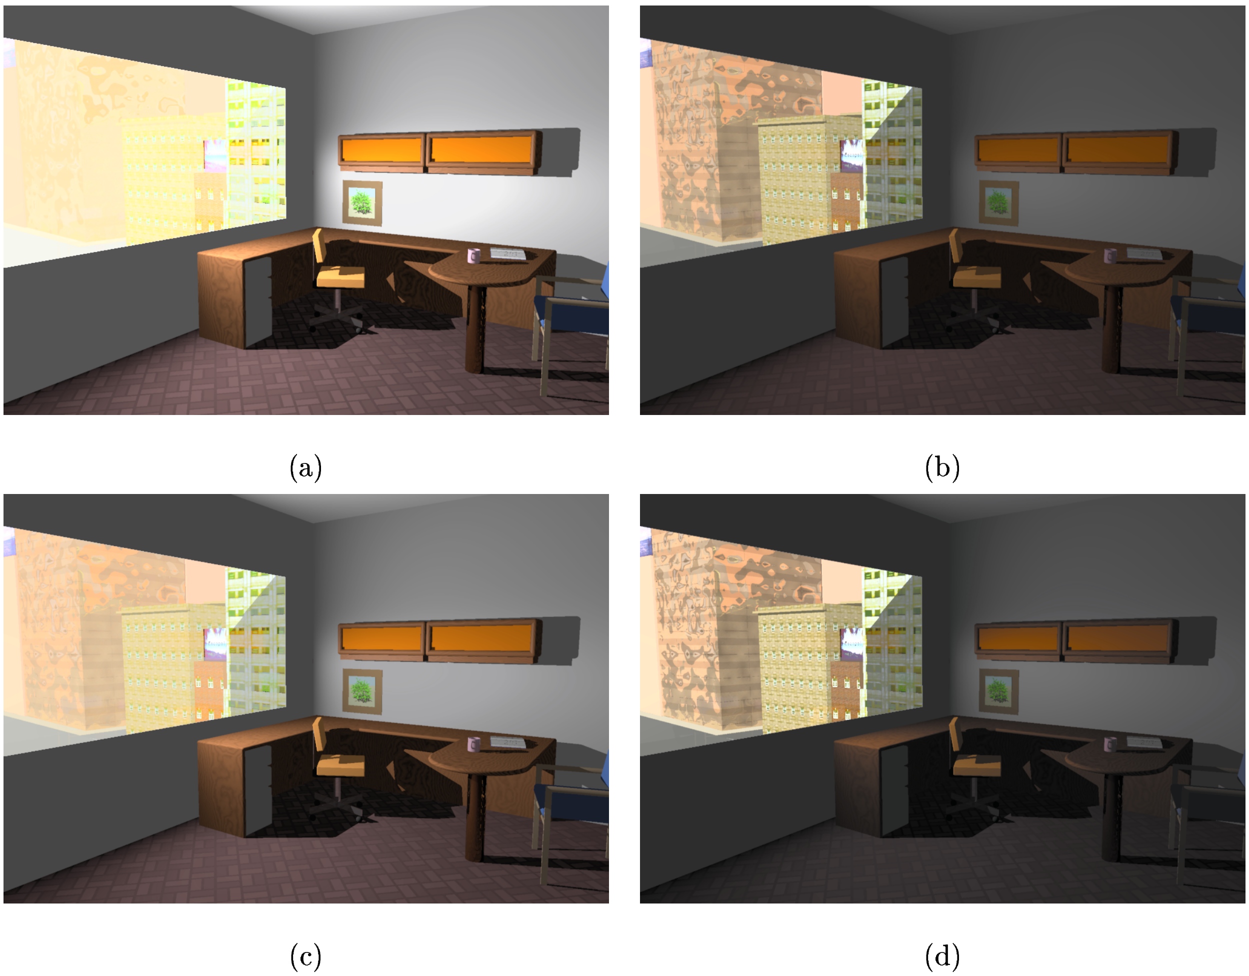

Figure 10: For each region we display the image with maximum entropy in the region. Images that visualize (a) region 0, (b) region 1 and (c) image with maximum global entropy. (d) Image obtained with Larson's algorithm

5 Conclusions

In this paper, we have presented a new paradigm for the reproduction and visualization of information in high dynamic range images. We argued for the use of a set of images instead of a single one as in traditional approaches.

More than being the last word about the problem of visualizing high dynamic range data, with this work we attempted to illustrate the intrinsic limitation of working with only one image. We showed how going for more than one image we could obtain a simple and nice solution to the problem of complete information visualization of high dynamic range images.

A number of questions remain open after this work. The specific algorithm here described for the computation of the set of images is just a particular example, and others should be developed. One of the crucial additional points is how to find the minimal number of images required to visualize all the relevant information. These images have also to be pleasant and hopefully with smooth transitions among them. We hope that the work here presented will open the door to works on these and other relevant questions.

A Larson's algorithm

For completeness, we now briefly describe the histogram adjustment used in Larson's algorithm for tone reproduction [5].

The algorithm adjusts the histogram based on a human contrast sensitivity function [4]:

$$\Delta L(L_a) = \begin{cases} -2.86 & \log_{10}(L_a) < -3.94 \\ (0.405\log_{10}(L_a) + 1.6)^{2.18} - 2.86 & -3.94 \leq \log_{10}(L_a) < -1.44 \\ \log_{10}(L_a) - 0.395 & -1.44 \leq \log_{10}(L_a) < -0.0184 \\ (0.249\log_{10}(L_a) + 0.65)^{2.7} - 0.72 & -0.0184 \leq \log_{10}(L_a) < 1.9 \\ \log_{10}(L_a) - 1.255 & \log_{10}(L_a) \geq 1.9 \end{cases}$$

We use the following notation: L_w is the world luminance, $B_w = \log(L_w)$ is the world brightness, L_d is the display luminance, B_d is the display brightness, L_{dmin} and L_{dmax} are the minimum and maximum display luminance, N is the number of luminance histogram bins, T the total number of samples, $f(b)$ is the frequency count at bin b, $P(b)$ is the probability of bin b, and Δb is the bin size.

To guarantee that the contrast in the output image does not exceed the contrast in the original scene, the derivative of the mapping is limited to:

$$\frac{dL_d}{dL_w} \leq \frac{\Delta L(L_d)}{\Delta L(L_w)}$$

Using the definition of the histogram we found that the number of points in each bin must satisfy

$$f(B_w) \leq \frac{\Delta L(L_d)}{\Delta L(L_w)} \frac{T \Delta b L_w}{[\log(L_{dmax}) - \log(L_{dmin})] L_d}$$

To solve this equation we apply an iterative procedure until a tolerance criterion is fulfilled. After that, we have a modified histogram that enables us to map the input to output levels using:

$$B_d = \log(L_{dmin}) + [\log(L_{dmax}) - \log(L_{dmin})] P(B_w)$$

Furthermore, the algorithm takes into account some aspects of human visual perception, including limitations concerning glare, color and acuity.

Acknowledgments

The Memorial church image was obtained from Paul Debevec's web page, www.cs.berkeley.edu/debevec. The office and bathroom images were obtained from Larson's web page, positron.cs.berkeley.edu/gwlarson. The authors thank Jack Tumblin for explaining his work [13], and Marcelo Bertalmio for fruitful discussions. This work was partially supported by a grant from the Office of Naval Research ONR-N00014-97-1-0509, the Office of Naval Research Young Investigator Award, the Presidential Early Career Awards for Scientists and Engineers (PECASE), a National Science Foundation CAREER Award, by the National Science Foundation Learning and Intelligent Systems Program (LIS), and by CSIC - Universidad de la República, Uruguay.

References

[1] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *SIGGRAPH 97 Conference Proceedings*, pages 369–378, 1997.

[2] J. DiCarlo and B. Wandell. Rendering High Dynamic Range Images. In *Proceedings of the SPIE, Image Sensors.*, volume 3965, January 2000.

[3] J.-O. Eklundh and A. Rosenfeld. Peak detection using difference operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1(3):317–325, July 1979.

[4] J. Ferwerda, S. Pattanaik, P. Shirley, and D. Greenberg. A Model Of Visual Adaptation for Realistic Image Synthesis. In *SIGGRAPH 96 Conference Proceedings*, pages 249–258, 1996.

[5] G. W. Larson, H. Rushmeier, and C. Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October-December 1997.

[6] S. Nayar and T. Mitsunaga. High Dynamic Range Imaging: Spatially Varying Pixel Exposures. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 472–479, June 2000.

[7] S. Pattanaik, J. Ferwerda, M. Fairchild, and D. Greenberg. A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display. In *SIGGRAPH 98 Conference Proceedings*, pages 287–298, 1998.

[8] radsite.lbl.gov/radiance/HOME.html. *Radiance.*

[9] Y. Y. Schechner and S. K. Nayar. Generalized Mosaicing. In *IEEE International Conference Computer Vision*, pages 369–378, 2001.

[10] C. Schlick. Quantization Techniques for Visualization of High Dynamic Range Pictures. In *Photorealistic Rendering Techniques, Proceedings of the 5th Eurographic Rendering Workshop*, pages 7–20. Springer-Verlag, June 1994.

[11] J. Tumblin, J. Hodgins, and B. Guenter. Two Methods for Display of High Contrasted Images. *ACM Transactions on Graphics*, 18(1):56–94, January 1999.

[12] J. Tumblin and H. Rushmeier. Tone Reproduction for Realistic Images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

[13] J. Tumblin and G. Turk. LCIS: A Boundary Hierarchy for Detail-Preserving Contrast Reduction. In *SIGGRAPH 99 Conference Proceedings*, 1999.

[14] B. Wandell, P. Catrysse, J. DiCarlo, D. Yang, and A. E. Gamal. Multiple Capture Single Image Architecture with a CMOS Sensor. In *Proceedings of the International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives. Japan.*, pages 11–17, 1999.