

Volume Rendering for High Dynamic Range Displays

Abhijeet Ghosh, Matthew Trentacoste, Wolfgang Heidrich
The University of British Columbia*



Figure 1: Screen photographs of volume rendering of the CT head dataset on a HDR display, captured at 2 f-stops apart, illustrating the representable dynamic range.

ABSTRACT

Dynamic range restrictions of conventional displays limit the amount of detail that can be represented in volume rendering applications. However, high dynamic range displays with contrast ratios larger than 50,000 : 1 have recently been developed. We explore how these increased capabilities can be exploited for common volume rendering algorithms such as direct volume rendering and maximum projection rendering. In particular, we discuss distribution of intensities across the range of the display contrast and a mapping of the transfer function to a perceptually linear space over the range of intensities that the display can produce. This allows us to reserve several just noticeable difference steps of intensities for spatial context apart from clearly depicting the main regions of interest. We also propose generating automatic transfer functions for order independent operators through histogram-equalization of data in perceptually linear space.

Keywords: Hardware – Novel Display Technologies; Rendering – Perceptually Based Rendering, Volume Rendering, Transfer Functions, Just Noticeable Differences (JNDs)

CR Categories: B.4.2 [INPUT/OUTPUT AND DATA COMMUNICATIONS]: Input/Output Devices—Image display; I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Display algorithms; I.4.0 [IMAGE PROCESSING AND COMPUTER VISION]: General—Image displays; I.4.10 [COMPUTER GRAPHICS]: Image Representation—Volumetric.

1 INTRODUCTION

Direct volume rendering has proven extremely useful for the visualization of medical and scientific data sets. One of its advantages is that transfer functions can be used to segment out interesting parts of the volume, while in principle keeping other information present to provide context useful for navigation.

Unfortunately, the low dynamic range of conventional displays limits the usefulness of this approach: for optimal contrast in the regions of interest, one has to adjust the transfer function such that most of the available intensity and opacity levels are used for very specific density values. Consequently, very little precision remains

for other density values to provide spatial context. The non-linear gamma curve of such display devices helps, but the problem remains as there are not enough addressable intensity values.

Recently developed display technology with a much higher dynamic range promises to be useful for solving this problem. Seetzen et al. [17] describe two such systems, one with projector-based illumination at a contrast ratio of 50,000 : 1 and a peak intensity of $2,700\text{cd}/\text{m}^2$, and one with LED illumination reaching a contrast ratio of $> 150,000 : 1$ and a peak intensity of $8,500\text{cd}/\text{m}^2$. Typical desktop displays have a contrast of about 400 : 1 with a maximum intensity of $300\text{cd}/\text{m}^2$, although special medical displays can perform a factor of 2-3 times better.

In this paper, we investigate the use of this HDR technology for volume rendering (Figure 1). In particular, we describe the use of transfer functions in perceptually linear space over the range of intensities representable by the display. We also describe adaptations of transfer functions to better represent spatial context and automatic generation of transfer functions based on JND space histogram-equalization. Finally, we explore related techniques for order independent volume rendering algorithms such as maximum intensity projection and summation, both of which are useful for x-ray style rendering.

The remainder of this paper is organized as follows: Section 2 briefly recounts previous work in volume rendering, HDR display technology, and the aspects of human visual perception that are relevant to our work. Section 3 discusses methods for deriving transfer function that yield in a perceptually linear usage of the contrast range, while Section 4 explores the possibilities for automatically adapting user-defined transfer functions to provide spatial context for navigation.

2 RELATED WORK

2.1 Direct Volume Rendering

A lot of recent work has focused on deriving transfer functions automatically from volume data. Much of this work analyzes the histogram of the volume densities, sometimes combined with gradients [6] or curvature [7]. Several different user interfaces have been proposed, namely [10, 9, 8].

Very recently, Mora and Ebert have argued that traditional direct volume rendering may not be the best way to visualize volume data

*e-mail: {ghosh, mmt, heidrich}@cs.ubc.ca

since important features may be occluded by less important parts of the volume. They propose *order independent* volume rendering, a framework that generalizes both maximum intensity projection and summation based methods, which roughly correspond to x-ray style rendering. They propose stereo as a way to compensate for the loss of depth cues arising from order independent methods, and argue that stereo is, in fact, more effective for order independent methods than for direct volume rendering.

Independent of a specific volume rendering method, it is important to evenly distribute information across the range of contrasts that can be shown on a given display. A first step in this direction was recently taken by Potts and Möller [15], who proposed to specify transfer functions for direct volume rendering on a logarithmic scale. However, for modern display devices (especially the high dynamic range displays that have recently evolved), it is also necessary to take into account the limitations of human contrast perception for various intensity levels. In this paper we analyze the intensities generated on the screen by various rendering algorithms, and propose to adapt transfer functions to take these perceptual effects into account. The goal is therefore to optimize the *perceptible* contrast generated in the final image.

2.2 HDR Displays

Conventional desktop display systems such as CRTs or LCD panels have dynamic ranges of about 400 : 1 and a maximum intensity of about $300cd/m^2$. In recent work, Seetzen et al. [17] describe two setups that combine conventional low dynamic range display technology to form a high dynamic range display. In the first setup, a video projector replaces the backlight of a conventional LCD panel. This way, the light arriving from the projector is filtered by the semi-transparent LCD panel. Seetzen et al. measured a dynamic range of about 50,000 : 1 with a maximum intensity of 2,700 cd/m^2 for this setup.

The same authors also developed a second system in which the projector is replaced with a low-resolution LED array. This is possible since the *local* contrast that the human eye can perceive is limited. This second setup achieves a top intensity of 8,500 cd/m^2 with a contrast of over 150,000 : 1.

Seetzen et al. describe the image processing operations necessary to factorize floating point images (representing absolute luminances) to drive the front panel and the back lighting of the HDR displays. These operations can be implemented on GPUs for integration into interactive rendering systems. In our work, we use the same algorithms as a backend for our volume renderer.

2.3 Human Perception and HDR Imaging

In recent experiments, Muka and Reiker [13] have determined that over the dynamic range of conventional displays the perceptual difference between an 8-bit digital display and a 10-bit or higher bit depth is minimal, and in some cases even non-existent. From this result, we can conclude that HDR display technology such as the one described in Section 2.2 is essential for displaying more visually distinct intensity levels.

Over the range of illumination levels representable by the HDR displays mentioned above, the sensitivity of the human visual system is highly non-linear: at low luminance levels, smaller differences are perceivable than at high luminance levels. This property is formally described by the notion of just noticeable differences (JND). One JND is the smallest detectable intensity difference for a given illumination level. For the intensity range covered by the HDR display technology mentioned above, Barten [1, 2] has derived a psychophysically validated model to characterize JNDs. Based on this work, an analytical function for computing JNDs was included in the DICOM Grayscale Standard Display Function [3].

A plot of the JND curve over the relevant intensity range is shown in Figure 2.

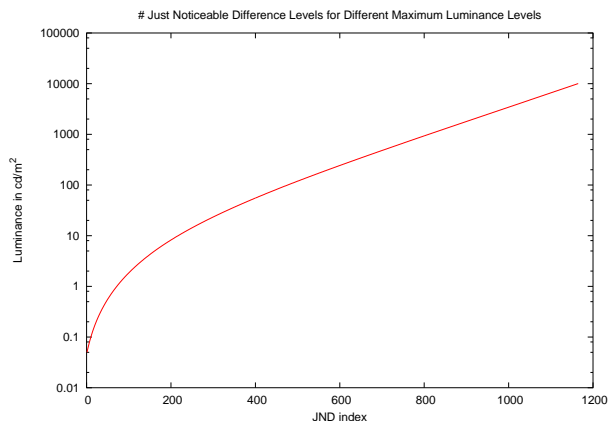


Figure 2: The number of just noticeable difference (JND) steps for different maximum intensities according to the DICOM standard.

For the projector-based HDR display of Seetzen et al. [17], this model predicts 962 JND steps, while for the LED-based display it predicts 1139 JND steps. Both display technologies can produce intensities at a finer granularity, but the human visual system cannot discriminate between those. To make optimal use of the contrast of a display, the intensities produced by an algorithm therefore must be linear in JND space, *not* the physical space. This is the focus of our work.

3 TRANSFER FUNCTIONS FOR HDR DISPLAYS

In the following we describe how to perceptually optimize user specified transfer functions for the different rendering algorithms. We start by discussing the case of direct volume rendering, and then move on to order independent methods (summation and maximum intensity projection).

3.1 Direct Volume Rendering

To derive an approximately perceptually linear formulation of the transfer function for direct volume rendering, consider the emission absorption volume rendering equation in the notation of Sabella [16]:

$$I(a, b) = \int_a^b C\rho(u) \cdot e^{-\int_a^u \tau\rho(t)dt} du \quad (1)$$

Here, ρ is the volume density at point u , C is an emission constant (i.e. $C \cdot \rho$ is the emitted energy per unit length), and τ is the absorption constant (i.e. $\tau \cdot \rho$ is the absorption per unit length).

If we assume that the density is constant over a segment of the integral, we get

$$I(a, b) = \frac{C\alpha}{\tau} \quad \text{with} \quad \alpha := 1 - e^{-\tau \int_a^u \rho(t)dt}, \quad (2)$$

as described by Max et al.[11]. In other words, the transparency α of the integral varies exponentially with the volume density. For this reason, Potts and Möller [15] argue that the transfer function that is used to derive densities from volume data values should be specified on a logarithmic scale.

However, as described in Section 2.3, the intensities themselves are not perceived linearly by the human observer. In order to make optimal use of the intensity range delivered by the HDR displays,

the just noticeable differences have to be taken into account. In particular, let f_{JND} be the JND function from the Dicom standard [3], and f_{JND}^{-1} be its inverse (i.e. the function mapping intensity values to just noticeable differences). The perceived intensity level in JND space is then:

$$I_{perceived}(a, b) = f_{JND}^{-1}(I(a, b)) = f_{JND}^{-1}\left(\frac{C\alpha}{\tau}\right). \quad (3)$$

The densities ρ should therefore be specified as

$$\rho := \log(f_{JND}^{-1}(\rho')), \quad (4)$$

where ρ' , the original volume densities, are now mapped approximately linearly to just noticeable intensity differences in the final image (Figure 3).

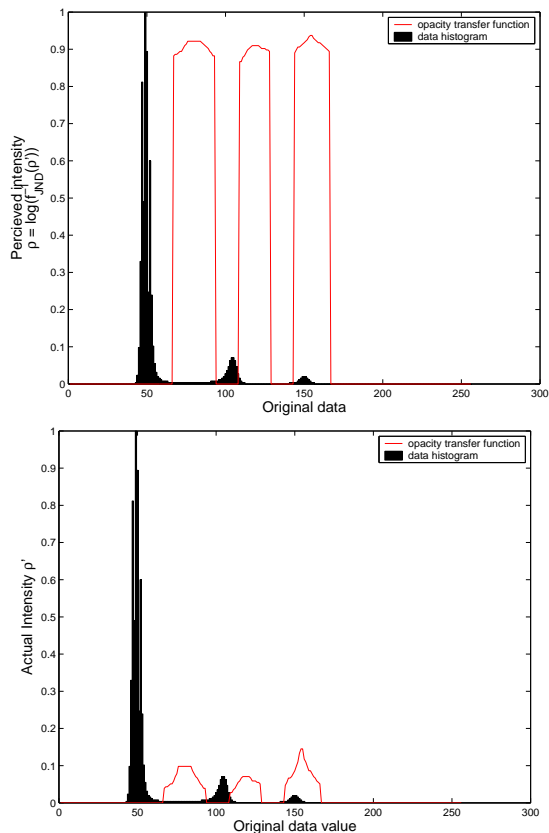


Figure 3: Perceptually linear transfer function specification for the tooth dataset (Figure 6). Top: perceived intensity levels ρ specified in JND space. Bottom: actual intensity ρ' used in volume rendering.

Of course the volume densities are not constant along the ray in practice. As a consequence, if a segment of the transfer function is changed, but others remain the same, then the actual intensity change on the screen can still be non-linear. However, we find that the correspondence between transfer function and response of the display is much more direct and easy to control if we use the mapping described above (see Section 5).

3.1.1 Color

The discussion so far has only considered just noticeable differences in intensity, and has ignored color. In practical volume rendering applications, color is, however, an important means of vi-

ually segmenting volume data sets into different parts. Unfortunately, knowledge about human perception in environments with both color and high contrast is at this point limited: most of the perception experiments dealing with color differences were performed in low contrast settings, while the intensity JND work is based on monochromatic experiments.

In the absence of more perceptual models that take both color and high contrast into account, the best we can do at the moment is to treat the two aspects as being independent. This can be achieved by using a color space that separates intensity from chromaticity, for example the $L^*a^*b^*$ space. The user can then specify two transfer functions for the chrominance channels (a and b), while the luminance is computed using the algorithm introduced for the monochromatic case.

3.2 Order Independent Operators

A similar analysis can be performed for the order independent operators. In the case of the summation operator, it is easy to see that a change in the transfer function contributes linearly to a change in pixel intensity. For maximum intensity projection, a change only occurs if the value changed is actually the maximum along a given viewing ray. If this is the case, however, then the effect is again linear. In both cases, the transfer function should therefore only be modified by the inverse of the JND curve.

In addition to this simple perceptual adjustments, order independent methods are also amenable to more sophisticated methods for automatically generating transfer functions. In practical volume rendering applications, a linear ramp transfer function is often used as a starting point for exploration. Mora and Ebert [12] showed that this can be a reasonable choice for order independent methods, although it does not tend to work very well for direct volume rendering.

Several researchers have focused on deriving transfer functions automatically from volume data [14]. The primary focus has been on analysis of the distribution of histogram values and sometimes combined with other features in the volume such as gradients [6] [4] or curvature [7].

We propose to perform histogram equalization on the JND-corrected volume data to generate the intensity transfer function. This equalization is done by first constructing the normalized cumulative histogram. We then normalize the histogram of the data values such that the intensity distribution is uniform in JND space. This allows us to generate a perceptually linear transfer function which maps most of the interesting data into the visible range (Figure 4). The generated transfer function can also be used as an intuitive starting point to explore the data and segment it in a more customized manner.

We find that this method works well for the summation operator. For the maximum operator, however, the histogram of the volume densities is not always a good predictor for the histogram of the pixel values, since the maximum operator is strongly view-dependent. For this reason, we also consider performing the histogram equalization in image space after projection. This is easy to implement for the maximum operator since it requires only one unnormalized value to be stored per pixel. In Section 5 we show several examples using the six order independent operators proposed by Mora and Ebert [12]: summation and maximum applied to original data values, to their gradient magnitude and to the *product* of data value and gradient magnitude.

Unfortunately, the theoretical motivation behind JND space histogram equalization does not apply to volume rendering. This is because the perceptually linear mapping of the data goes through another non-linear mapping in the form of exponential fall-off due to the volume rendering integral (Equation 1). In practice, however, we find that the method can often still be used to provide a good

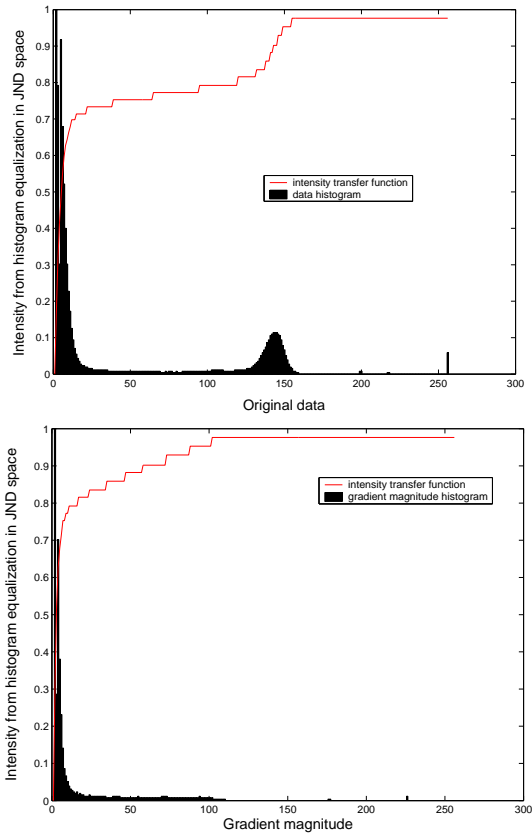


Figure 4: Automatic transfer function generation based on JND space histogram equalization for the CT engine dataset. Top: transfer function based on equalization of data density values. Bottom: transfer function based on equalization of gradient magnitude (as used in Figures 8 and 9).

starting point for a transfer function. We provide several examples of results for direct volume rendering in Section 5.

4 AUTOMATICALLY PROVIDING SPATIAL CONTEXT

Something we might want to do on a HDR display is to reserve a small range of density values for parts of the volume that are not directly the focus of attention, but can provide context for navigation within the volume. A perceptually linear space for specifying transfer functions is useful in this regard, since it allows us to directly estimate what contrast range will be used for this purpose.

For example, on the projector-based HDR display [17], we can set aside a portion of the 962 displayable JND steps for providing spatial context in this form. In our implementation, we use information from the data itself, in the form of gradient magnitude, to highlight context. Volume areas of high gradient magnitude correspond to distinctive isosurfaces that can be rendered dimly to provide context for navigation. We found that setting aside about 400 JNDs for providing this context is a good tradeoff for the projector based HDR Displays. We expect that the best choice will depend on the contrast and intensity range of the individual display used. For our display, this choice leaves more than 560 JNDs for depicting the data values the user is currently interested in, which is still more than twice the precision of conventional displays.

Given a user provided transfer function that segments out the object of interest and sets the densities of all other regions to 0, we can adapt this transfer function in the following way: all non-

zero entries in the old transfer function are linearly mapped to JND steps reserved for focus (in our case 400 . . . 962). The regions that were zeroed out previously are replaced by piecewise linear segments mapping gradient magnitude values to the range 0 . . . 400 (Figure 5).

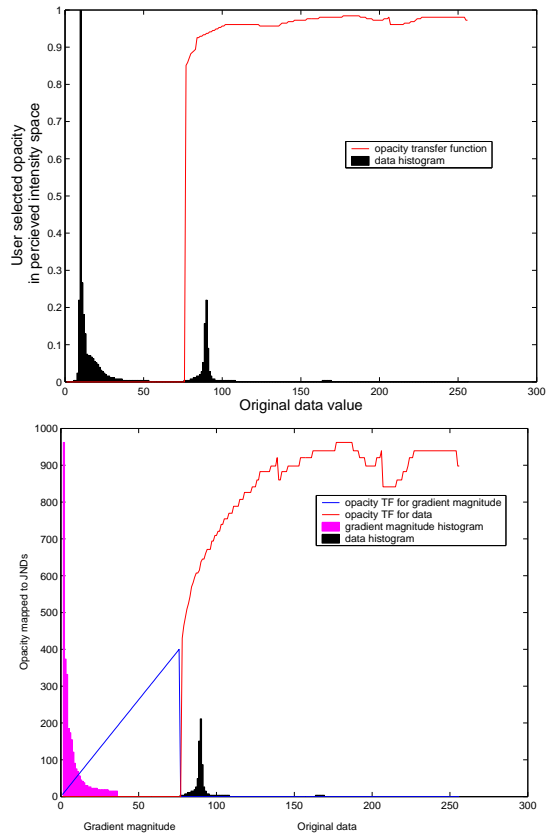


Figure 5: Automatic context generation for the CT head dataset (Figure 7). Top: User specified opacity transfer function for data (yellow histogram). Bottom: Automatic context for unselected data in the form of gradient magnitude (orange histogram) being mapped to 0 . . . 400 JNDs, while the selected data is mapped to 400 . . . 962 JNDs reserved for focus.

5 RESULTS AND DISCUSSION

We use the perceptually linear mapping described in Section 3.1 to specify the opacity transfer function for various medical and scientific datasets. We present results on the projector based HDR display as well as comparisons with tone mapped versions on a regular display device. Here, we used a log-linear tone-mapper with gamma correction as implemented in HDRShop [5]. The CT tooth dataset has many distinct isosurfaces very close together in intensity space and a logarithmic transfer function (Figure 3) significantly aids in isolating these isosurfaces (Figure 6, left). Note that the HDR display clearly shows a lot more detail in the tooth than the tone-mapped version on the regular display. The same non-linear mapping was used to isolate the sinuses and a thin layer of skin around the skull in the CT head dataset (Figure 6, right). Again note the details in the eye sockets and the skull surface as shown on the HDR display compared to the version shown on the regular display device. The HDR sequences also convey a better sense of relative depth for various features. In this case, the tooth and the head datasets were both rendered with lighting and shading.




Figure 6: Screen photographs of volume rendering of the CT tooth and the CT head datasets. The left-right image pairs were captured at 4 f-stops apart. Top row: As displayed on the HDR display; Bottom row: tone-mapped version displayed on a regular LCD panel.

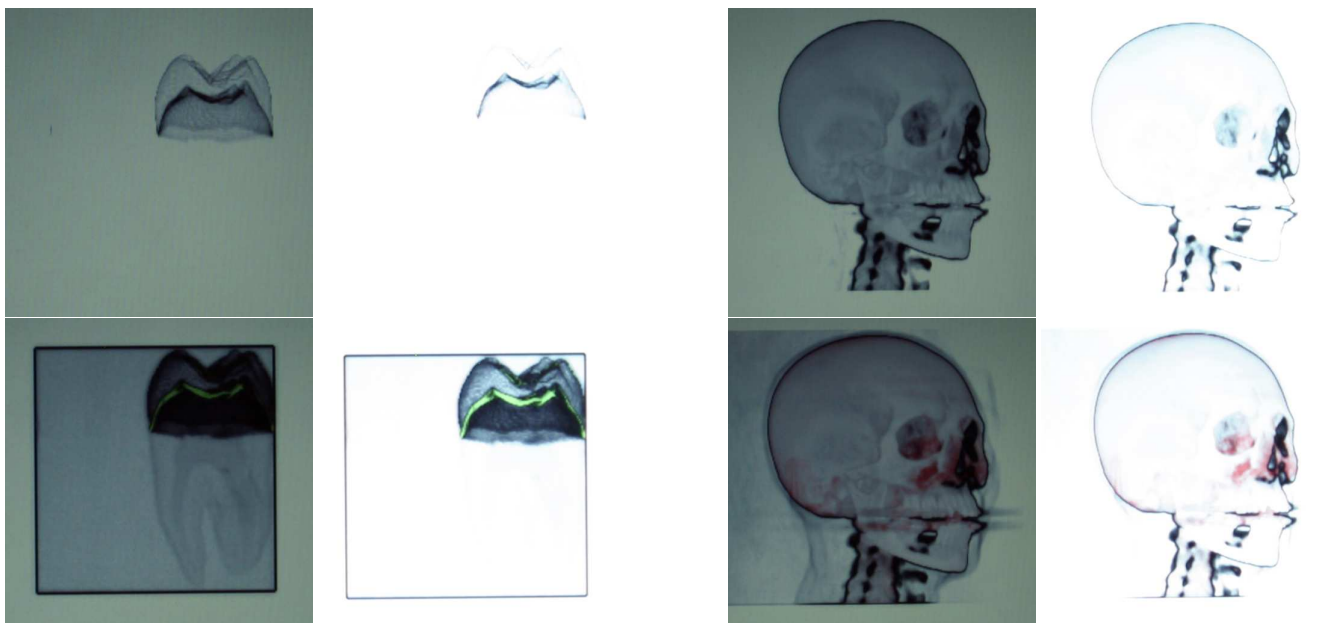


Figure 7: Automatic context generation for both the CT tooth and the CT head datasets by adapting the transfer function and remapping the intensities in JND space. The left-right image pairs were captured at 4 f-stops apart. Top row: User selected focus; Bottom row: Automatically generated context.




Figure 8: JND space histogram equalization for generating automatic transfer functions for the *sum* operator. The left-right image pairs were captured at 4 f-stops apart. Top: transfer function generation by histogram equalization; Bottom: JND space *linear ramp* transfer function. Left-pairs: CT engine dataset and CT head using gradient magnitude; Right-pairs: CT engine dataset and CT tooth using *product*.

Reserving JNDs for context as described in Section 4 can be very useful as a semi-automatic way of generating transfer functions. We use the gradient magnitude of the data for visualizing context since gradient magnitude defines object boundaries. We selected a threshold of 0 . . . 400 JNDs with the projector-based display for visualizing context as this nicely separated out the various isosurfaces for most of the datasets. This threshold can also be set in a data-driven way, for example by examining the histogram of the gradient magnitude. The context generated in this way for the CT tooth and head datasets (Figure 7, left and 7, right) results in semi-automatic isolation of interesting isosurfaces, such as the sinuses, ear and skin in the CT head dataset and the roots in the tooth, similar to that obtained from full user selection in Section 3.1. Also note how the context is mostly saturated in the long exposure shots as it occupies only a small portion of the intensity space in order to retain details in the focus. The tooth and head datasets were rendered without lighting and shading in this case to clearly illustrate the effect of

JND space context generation. Note that the JND space mapping for focus and context was applied only to the opacity transfer function and color was manually assigned.

Our proposed JND space histogram equalization in Section 3.2 provides a way to automatically generate a perceptually linear transfer function which maps most of the interesting data into the visible range. It also serves as an improved starting point for further exploration of the data. As pointed out previously, this is directly applicable for automatically generating transfer functions for order independent operators in volume rendering. The histogram equalization is applied in data space for the summation operator and in image space for the maximum operator. We apply it to the six order independent models [12]: summation and maximum applied to original data values, to the gradient magnitude and to the product of the data value and gradient magnitude. We present comparisons of renderings with automatically generated transfer functions through histogram equalization in JND space with those using a linear ramp




Figure 9: JND space histogram equalization for generating automatic transfer functions for the *max* operator. The left-right image pairs were captured at 4 f-stops apart. Top: transfer function generation by histogram equalization; Bottom: JND space *linear ramp* transfer function. Top-left: CT engine dataset using gradient magnitude; Top-right: CT engine dataset using *product*; Bottom-left: CT visible human male dataset using data density; Bottom-right: CT visible human male dataset using *product*.

in *JND space* for these operators.

With summation, the operator applied to gradient magnitude and the product provides the most compelling visualizations as both cases highlight the distinct isosurfaces in the data. Figure 8 presents both results for the CT engine dataset, as well as gradient magnitude image for the CT head, and product for the CT tooth.

Figure 9 presents the results of applying the maximum operator to both gradient magnitude and the product for the CT engine dataset, as well as to original data and the product for the CT Visible Human Male dataset. Note that here we compare results of JND space histogram equalization in image space to a linear ramp in JND space. In the case of both summation and maximum intensity projection, the JND space equalization leads to a better visualization of various features in the volume.

6 CONCLUSIONS

One of the main advantages of direct volume rendering for visualization of medical and scientific datasets is the usage of transfer functions in order to segment out interesting parts of the volume, while in principle keeping other information present to provide context useful for navigation.

Conventional displays with low dynamic range provide very little precision for spatial context as most of the precision is used up for visualizing specific segmented regions of the volume. In this paper, the use of recently developed HDR display technology for volume rendering is investigated. In particular, we examine the creation of transfer functions in perceptually linear space over the range of intensities representable by the display. This is done by mapping intensity values to just noticeable differences and defining transfer function in JND space. Also described is the adaptation

of transfer functions to better represent spatial context by reserving JND levels for both focus as well as context. Automatic transfer function generation in perceptually linear space is also presented through histogram equalization in the JND space.

There are a number of opportunities for future work. In particular, the treatment of color is presently based on the assumption of independence between just noticeable differences in luminance and chrominance. The availability of the HDR display technology would now allow the design of perceptual experiments to verify this assumption, or to derive better models, which could, in turn, be used to improve the methods presented here.

Our current analysis does not account for the changes of the intensities by lighting and shading computations. If shading is to be included into the contrast optimization, one would probably not want to give the same priority to these shading-based differences as to differences based on actual data values. A possible compromise could be to limit the influence of the shading operations to a small number of JNDs, similar to the way we currently create navigational context.

7 ACKNOWLEDGMENTS

The first author was supported by an ATI Technologies Fellowship. We would like to thank Helge Seetzen and Sunnybrook Technologies for providing the HDR display as well as technical support and feedback.

REFERENCES

- [1] P.G.J. Barten. Physical model for the contrast sensitivity of the human eye. In *Proc. SPIE*, volume 1666, pages 57–72, 1992.
- [2] P.G.J. Barten. Spatio-temporal model for the contrast sensitivity of the human eye and its temporal aspects. In *Proc. SPIE*, volume 1913-01, 1993.
- [3] DICOM. *Digital Imaging and Communications in Medicine (DICOM)*, chapter Part 14: Grayscale Standard Display Function. <http://www.medical.nema.org/dicom/2001.html>, 2001.
- [4] D. Ebert, C. Morris, P. Rheingans, and T. Yoo. Designing effective transfer functions for volume rendering from photographic volumes. *IEEE Trans. on Visualization and Computer Graphics*, 8(2):183–197, April-June 2002.
- [5] HDRShop. <http://www.debevec.org/hdrshop/>.
- [6] G. Kindlmann and J. Durkin. Semi-automatic generation of transfer functions for direct volume rendering: Methods and applications. In *Proc. of Visualization '98*, pages 79–86, 1998.
- [7] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Möller. Curvature-based transfer functions for direct volume rendering. In *Proc. of Visualization '03*, 2003.
- [8] J. Kniss, G. Kindlmann, and C. Hansen. Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):270–285, 2002.
- [9] A. König and E. Gröller. Mastering transfer function specification by using VolumePro technology. In *Proc. of Spring Conference on Computer Graphics (SCCG) '01*, pages 279–286, April 2001.
- [10] J. Marks, B. Andalman, P. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proc. of SIGGRAPH '97*, pages 389–400, 1997.
- [11] N. Max, P. Hanrahan, and R. Crawfis. Area and volume coherence for efficient visualization of 3D scalar functions. In *Proc. of Volume Visualization '90*, pages 27–33, November 1990.
- [12] B. Mora and D. Ebert. Instant volumetric understanding with order-independent volume rendering. *Computer Graphics Forum*, 23(3):489–497, 2004.
- [13] E. Muka and G.G. Reiker. Reconsidering bit depth for radiological images – is eight enough? In *Proc. SPIE*, volume 4686, pages 177–188, 2002.
- [14] H. Pfister, B. Lorensen, C. Bajaj, G. Kindlmann, W. Schroeder, L. Avila, K. Martin, R. Machiraju, and J. Lee. The transfer function bakeoff. *IEEE Computer Graphics and Applications*, 21(3):16–22, May-June 2001.
- [15] S. Potts and T. Möller. Transfer functions on a logarithmic scale. *Proc. of Graphics Interface*, pages 57–63, May 2004.
- [16] P. Sabella. A rendering algorithm for visualizing 3D scalar fields. In *Proc. of SIGGRAPH '88*, pages 51–58, August 1988.
- [17] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs. High dynamic range display systems. *ACM Transactions on Graphics (Proc. of Siggraph)*, 23(3):760–768, 2004.