

Volumetric High Dynamic Range Windowing for Better Data Representation

Dirk Bartz, Benjamin Schnaidt, Jirko Cernik,
Ludwig Gauckler, Jan Fischer, Angel del Rfo

WSI-2005-03
January 2005

Visual Computing for Medicine
Graphisch-Interaktive Systeme
Wilhelm-Schickard-Institut
Universität Tübingen
D-72076 Tübingen, Germany
e-mail: bartz@gris.uni-tuebingen.de
WWW: <http://www.gris.uni-tuebingen.de/areas/vcm>

© WSI 2005
ISSN 0946-3852

ABSTRACT

Volume data is usually generated by measuring devices (eg. CT scanners, MRI scanners), mathematical functions (eg., Marschner/Lobb function), or by simulations. While all these sources typically generate 12bit integer or floating point representations, commonly used displays are only capable of handling 8bit gray or color levels. In a typical medical scenario, a 3D scanner will generate a 12bit dataset, which will be downsampled to an 8bit per-voxel accuracy. This downsampling is usually achieved by a linear windowing operation, which maps the active full accuracy data range of 0 up to 4095 into the interval between 0 and 255.

In this paper, we propose a novel windowing operation that is based on methods from high dynamic range image mapping. With this method, the contrast of mapped 8bit volume datasets is significantly enhanced, in particular if the imaging modality allows for a high tissue differentiation (eg., MRI). Henceforth, it also allows better and easier segmentation and classification. We demonstrate the improved contrast with different error metrics and a perception-driven image difference to indicate differences between three different high dynamic range operators.

CR Categories: I.4.3 [Enhancement]: Filtering, Gray-scale manipulation, Sharpening and deblurring; I.4.7 [Feature Measurement]: Feature representation; I.4.8 [Scene Analysis]: Intensity;

keywords: Volume Data, Non-Linear Data Mapping, High Dynamic Range Mapping, Windowing

1 INTRODUCTION

There are numerous sources for volumetric data. Among the most frequent ones are measuring devices – such as CT scanners, MRI scanners, etc. –, simulations of natural phenomena, or voxelization and implicit function schemes. Almost all volumetric data sources have in common that they generate voxel data with a float or integer data type. Very rarely, the original data is only an one-byte integer, the common data format of many volume rendering systems. Furthermore, most displays (gray-level or color), as they are used with virtually every computer, provide only an 8bit data fidelity per color channel. Therefore, the float and integer data types need to be mapped into the smaller data range, if only one color channel (eg., for gray-level representations) is used. This is in particular true in most medical applications, where a high dynamic range display is often prohibitively expensive to be used at every workplace. Up to now, the common remedy is a linear windowing operation, which maps a usually interactively selected data range into 8bits. This, however, means that the data has been downsampled to fit into 8bits, or lower and higher data values are clamped. Usually, the data is windowed by a combination of downsampling and clamped, based on standardized settings for specific imaging protocols.

In this paper, we present a new approach for a volume data windowing operator that employs techniques from high dynamic range imaging to map volume data from the original voxel size, usually 12bits¹, into a target voxel size, usually 8bit. 12bit voxels cannot

really be considered representing a truly high dynamic range. Nevertheless, we exploit the algorithms to enhance the contrast in the 8bit representation, but we do not require techniques that specifically address very high intensities. Note that this approach can also be used to increase contrast in native 8bit datasets (mapping an 8bit data range into a different 8bit data range), or in representations where the full data range is lost or not available for other reasons.

In contrast to the original image tone mapping operators, our operator examines the whole volume dataset and is therefore volumetric in nature. Although we have only explored datasets obtained by medical 3D scanners, we believe that this method is equally applicable to other scalar data types, such as single and double precision floats (of possibly truly high dynamic range), since the algorithm is already based on float numbers.

The operator is based on the Photographic Tone Reproduction operator (PTR) introduced by Reinhard et al. [16], but required several significant functional modifications for medical imaging. Furthermore, the volumetric nature of medical data required different processing strategies in order to cope with the increased space and time complexity.

In the remainder of the paper, we will discuss related work and introduce Reinhard’s PTR operator (Section 2). In Section 3, we describe our approach and its contributions to extend the PTR operator for volumetric data. Afterwards, we demonstrate the new windowing operator in Section 4. Finally, we draw conclusions and provide perspectives for future work.

2 RELATED WORK

Manual or automatic selection of an intensity window is a common step in the processing pipeline of medical visualization systems. Windowing is described in [5] and, in more detail, in [14]. Classical intensity windowing consists of selecting maximum and minimum intensities of interest. The values in this intensity range, which is usually defined within the 10 or 12bit pixel depth provided by typical medical scanners, are then mapped linearly to the gray-scale resolution of the display device. Thus the acquired medical images are normally converted to 8bit gray-scale images. Only expensive specialized display systems are capable of displaying images with higher pixel depths [4]. Consequently, intensity windowing is described as an integral part of a medical visualization system [12].

The linear mapping of pixel or voxel intensities from medical scanner measurements to 8bit gray values often does not produce satisfactory contrast. The intensities of certain types of tissue can be clustered within a very narrow range of values, making them difficult to distinguish on the display. The standard approach to improve contrast in images is histogram equalization [11], where the intensities of an image are transformed such that their distance in the intensity range depends on the respective weight in the image histogram. Although this non-linear mapping is based on the image data, it is a global operation and does not take local features into account.

More general mapping strategies and further improvements have been proposed in recent years, in order to maximize the quality obtained when displaying high dynamic range images. In computer graphics, this windowing process is usually known as *tone reproduction*, *tone mapping*, or *high dynamic range image mapping*.

Next to the global tone mapping operators (also known as *spatially uniform* or *single-scale*), like the previously referenced his-

¹3D medical image scanners usually generate data that requires up to 12bits per voxel. While the resulting data range of 0 to 4095 (translating

CT Hounsfield units into the positive range) is almost fully exploited by CT data, MRI data require only 9-10bits per voxel.

togram equalization, one can also find local operators (also known as *spatially varying* or *multi-scale*). While global operators perform the same mapping operation over the whole image, local operators can act distinctly over different regions in the image providing further control over the final result. Since the algorithm that we propose works on a multi-scale basis, we concentrate here on the most recent and relevant methods presented on local tone reproduction operators. A general survey of tone mapping techniques is given in [8].

Since the early work of Chiu et al. [7], the advantage of analyzing different parts of a high dynamic range image has been proven to establish a visually satisfactory tone mapping. Many different local operators for this purpose have been derived afterwards. They exploit the characteristics of the human visual system so that the necessary reduction of the dynamic range and its inherent detriment of accuracy lead to a minimum loss of visible quality. In [10] for instance, an attenuation of large luminance gradients is performed to compress these large gradients, while preserving fine details.

A different approach is presented in [9], which is based on the previous work by Tumblin et al. [18, 19]. In this case an edge-preserving filter (*bilateral filter*) is employed to decompose the images into two layers, a base layer and a detail layer. The contrast is reduced in the base layer, which contains the large-scale variations, while visibility is preserved in the detail layer.

Ashikhmin presented a multi-pass strategy for tone reproduction [3]. In a first stage, an estimation of the local adaptation luminance is performed at each point in the image. Then a compression function is applied to the values obtained during the first step in order to adapt to the required dynamic range of the display. In the third and final stage, relevant details are restored by a post-processing pass over the image. Each of these stages can be performed for a voxel at a time, thus the whole algorithm can be computationally combined into one pass.

Another tone mapping operator (PTR) has been proposed by Reinhard et al, in [16]. Their work is based on a classical method widely employed in photography to correctly reproduce real world high dynamic range scenes, called the Zone System [1]. In this technique, the scene is divided into print zones mapping the values from pure black to white. A subjective middle brightness region of the scene is taken and assigned to the middle gray print zone. Next, luminance readings are taken for light and dark regions to determine the dynamic range of the scene. It must be noted that the term dynamic range is used here in the photographic sense, that is, the ratio of the highest to the lowest luminance regions with visible details. This is not the typical computer graphics meaning as ratio of the highest to the lowest scene luminance. Using this method, it is assured that the maximum possible detail is retained, whenever an appropriate middle gray is found. Reinhard et al. propose a double pass strategy composed of an initial luminance mapping and a posterior automatic *dodging-and-burning* process. The initial mapping is a spatially uniform transformation where the tonal range of the output image is set in relation to the log-average luminance of the scene and a user-specified value for middle gray. During this mapping, high luminances are compressed. Then, in a second stage a local tone reproduction algorithm that applies *dodging-and-burning* is employed. This technique mimics the photographic *dodging-and-burning* process, where the portions of the print might receive a different exposure time from the negative, controlling the contrast on every region. In a similar way, contrast can be controlled locally over an entire region bounded by large contrast changes. These regions are determined using a center-surround function derived from the brightness perception model by Blommaert et al. [6].

3 VOLUMETRIC HIGH DYNAMIC RANGE WINDOWING

Until now, tone mapping has only been used for two-dimensional images. We present a straightforward, but novel variation of this method suitable to enhance a full volumetric dataset. Since we focus here on gray-level medical images, we will use the terms luminance and intensity interchangeably.

We modified the tone mapping method presented by Reinhard et al. [16] such that it provides a more comprehensible representation of radiological 2D images. In this course, we have implemented two variations of the original algorithm. One method performs high dynamic range windowing on individual slice images of the data volume. A second, significantly different approach works on the entire volume dataset. Since the volumetric variation performs computationally only somewhat more expensive than the slice-based version, we will only discuss the volumetric variation. The slice-based variation can also be seen as the volumetric variation with a constant z parameter. The visual differences of the volumetric and slice-based approach are shown in Figure 4e and f.

In Sections 3.1 and 3.2, we will give a summary of the modified (but largely similar) intensity mapping (which we call zone intensity mapping) and of the extended volumetric *dodging-and-burning* approach, based on the formulas presented by Reinhard et al. [16]. Note that while the theoretical basis of the algorithm largely follow Reinhard’s approach – we use a slightly different and easier notation –, our implementation varies significantly to achieve the necessary computational performance. Unless noted otherwise, most of the parameter settings are tuned the same way as suggested by Reinhard et al. in [16]. In particular a , s , α , and ε are specified identical, since these values produced good results for medical image data as well. Therefore, we will not specifically mention this fact in the following sections.

3.1 Zone Intensity Mapping

In a first step the intensity of an input voxel is scaled, so that a certain mean intensity is achieved. This scaled intensity is obtained by the computation of the log-average intensity \bar{I}_w as shown in Equation 1. Please note that not all volumetric datasets use the full 12bit data range ([0;4095]). In those cases, the mapping takes only the actively used bit range (eg., 11bits for [0;2047]).

$$\bar{I}_w = \exp\left(\frac{1}{N} \cdot \sum_{x,y,z} \log(\delta + I_w(x,y,z))\right) \quad (1)$$

In Equation 1, I_w denotes the intensity of the original voxel (world intensity). N is the total number of voxels, and δ is a small number used for ensuring that the logarithm in the sum can be computed for possible voxels with intensity 0. The log-average intensity \bar{I}_w is computed over all three-dimensions of the volume and represents the geometric mean of $\delta + I_w(x,y,z)$. In contrast to Reinhard [16], who focuses on images from photography, we focus on medical scanned data which consists of significantly larger numbers of black (zero) voxels that contain no or only little information. Using a δ close to zero like Reinhard would overemphasize these black voxels, hence we use $\delta = 1.0$ and compensate by subtracting $\delta = 1.0$ from \bar{I}_w , resulting in Equation 2:

$$\bar{I}_w = \exp\left(\frac{1}{N} \cdot \sum_{x,y,z} \log(1 + I_w(x,y,z))\right) - 1 \quad (2)$$

The voxel intensities are adapted using the calculated log-average intensity and a target key value a . The key value a determines the overall brightness of the resulting image. As suggested in the original article, we have found a key value of 0.18 to produce good results in most cases [16]. This key value is essentially motivated by the human visual system, which perceives 18% reflectance

Table 1: Computation time (s) of logarithms in Equation 2. An upper bound of 0 represents the log-computation for every voxel.

Dataset, Resolution	Time (s) for Upper Bound		
	0	10^{100}	10^{300}
CT Head, 512 x 512 x 324	18.15	4.38	4.12
CT Thorax, 512 x 512 x 168	9.49	2.28	2.14
CT Angiography, 512 x 512 x 194	10.81	2.61	2.44
MRI CISS, 512 x 512 x 54	3.01	0.72	0.69
CT Backpack, 512 x 512 x 373	18.93	4.84	4.31

as *middle gray*. This concept of a “normal key” works also for medical data, since it does usually not exceed 12bit data values (and hence can be considered neither high nor low key). Each of the voxels is scaled according to Equation 3.

$$I(x, y, z) = \frac{a}{I_w} \cdot I_w(x, y, z) \quad (3)$$

Using this adaptation, a comparable mean intensity is achieved for any given volume dataset. The scaled voxel intensity $I(x, y, z)$ is used as input for the following steps of the algorithm.

If we use only zone-intensity mapping (without the subsequent *dodging-and-burning*), we use Equation 4 with I_{MAX} as the maximum intensity found in the original input volume. This operator provides a general increased intensity and a slightly improved contrast.

$$I_c(x, y, z) = \frac{I(x, y, z) \left(1 + \frac{I(x, y, z)}{I_{MAX}}\right)}{1 + I(x, y, z)} \quad (4)$$

In terms of computational costs, the computation of the log values for the individual voxels is the single most expensive operation. Fortunately, Equation 2 can be reformulated as logarithm of the product of the argument of the logarithms: $\log(a \cdot b \cdot c) = \log(a) + \log(b) + \log(c)$. Albeit, the product over all voxels would quickly pass the accuracy bounds of 64bit floats, we can still combine a significant number of products, before we are getting close to this bound. We use 10^{100} as bound – which is still significantly below the upper bound of 64bit floats – and gain a speedup of four to five times compared to the direct implementation of Equation 2, as it can be seen in Table 1. It also shows that an increase to 10^{300} as upper bound does not significantly reduce the computing time.

3.2 Volumetric Dodging-and-Burning

The term *dodging-and-burning* was originally coined in photography and relates to a specific exposure technique that lightens or darkens regions in the final print [2]. While *dodging-and-burning* in photography is largely depending on the individual choice of the photographer, Reinhard’s *dodging-and-burning* is an automatic process. Its purpose is essentially to treat regions with similar contrast (a mostly homogeneous region) identical. In this process, the original algorithm tries to identify these regions by computing the convolution of the intensity with a Gaussian filter of a specific size, which is called the *scale* of the filter. The largest filter that still covers a region, which has a similar contrast, is selected for further processing.

We have devised a volumetric extension of this *dodging-and-burning* operator for the direct application in three-dimensions. Unlike the original algorithm, it does not work on a per-slice basis. Rather, both the filter kernels and the convolutions used for determining the correct scale are defined in 3D.

In contrast to global tone mapping, *dodging-and-burning* works on local regions. As mentioned above, in a first step, the size of this

local region is determined for each voxel. This size depends on the area around the voxel in which no large contrast changes occur.

In order to measure the contrast around a given voxel, we follow Reinhard et al., who used a center-surround function based on Blommaert’s model for brightness perception [6], which is based on spherical symmetric Gaussian filter kernels:

$$R_i(x, y, z, s) = \exp\left(-\frac{x^2 + y^2 + z^2}{(\alpha s^i)^2}\right) \quad (5)$$

Note that $R_i(x, y, z, s)$ is not yet normalized, which is done implicitly during processing later on. Also, we are using a slightly different notation than Reinhard et al. uses in Equations 5-8, which are, however, equivalent, since we iterate s instead of α .

For each of these kernels, the scale at which it is supposed to operate is given as s^i . Here s is the so-called center-surround ratio, which is the factor between two consecutive scales (Note that $s^0 = 1$, $s^1 = 1.6$, $s^2 = 1.6 \cdot 1.6$, etc.). The currently relevant scale is determined by the index i . We have chosen a fixed parameter value of $\alpha = 1/\sqrt{2} \approx 0.35$ in our experiments to ensure that the smallest scale essentially covers only the voxel itself. Although this value can be changed by the user, this fixed α generated satisfying results for medical datasets.

We can now obtain the three-dimensional convolution of the intensities of the volume data with these Gaussians, as shown in Equation 6, whereas we are using only $n \times n \times n$ filter kernel, where n is by default 5.

$$V_i(x, y, z, s) = I(x, y, z) \otimes R_i(x, y, z, s) \quad (6)$$

Since Gaussian kernels are separable, we can compute the convolution in three 1D passes, in contrast to one 3D pass [16, 13]. Thus, the computational complexity per-voxel is $O(n + n + n)$, in contrast to $O(n \times n \times n)$, resulting in only four-times of the computational costs for $n = 5$, instead of 48-times higher.

The computation of the 3D convolutions requires accessing voxel data from different slices. This significantly degrades memory access efficiency, resulting in considerably longer processing times. Unfortunately, not all participating slices fit into the L1-cache for fast processing. However, the separated convolution scheme as described above allows a much more compact memory representation. This way, the actual convolution fits easily into the L1-cache.

Please note that the frequency-based convolution – suggested by Reinhard et al. as convolution approach – has various other drawbacks compared to our truncated Gaussian filtering, since it requires the full (Fast) Fourier Transformation (FFT) of the data volume into frequency space. Furthermore, it would require this transformed data volume at several scales i , which would be prohibitive expensive for volumetric data, not only in computational costs, but also in memory requirements. Nevertheless, we tested this frequency space-based convolution and compared the voxel differences to our $n \times n \times n$ Gaussian filter convolution, with $n = 5, 11, 21$. The differences for $n = 5$ were largely negligible and for $n = 21$ not detectable. A similar result can be expected, if we compare the Gaussian filter matrix with a full-volume (full-image in 2D) sized filter. Therefore, we consider this matrix based convolution as an acceptable and significantly cheaper solution.

Now the center-surround function can be evaluated, which gives a measure for the activity (dynamics) at a given scale around the voxel (see Equation 7). Note that the center-surround ratio s is here already considered a (user specifiable) constant.

$$act_i(x, y, z) = \frac{V_{i-1}(x, y, z, s) - V_i(x, y, z, s)}{2^\phi a / (s^{i-1})^2 + V_{i-1}(x, y, z, s)} \quad (7)$$

The algorithm stops at the first scale index i for which $|act_i(x, y, z)|$ is larger than $\epsilon = 0.05$, and picks the previous scale

index $i - 1$. Choosing the respective convolution result V_{i-1} , a new intensity for the given voxel can be computed according to Equation 8. The operator specified in this equation overall provides a compression of high intensities. While this is not necessary for most medical image data, it can be used to compensate very high intensities caused by metal artifacts in CT data (eg., by metal tooth fillings in Fig. 1).

$$I_c(x, y, z) = \frac{I(x, y, z)}{1 + V_{i-1}(x, y, z, s)} \quad (8)$$

A further improvement of the intensity mapping in Equation 8 is achieved by adapting the computed output I_c to the maximum intensity (I_{MAX}) that is present in the volume (Equation 9), similar to Equation 4 of the zone intensity mapping.




Figure 1: Metal artifacts in CT Head dataset. (a) Linearly mapped slice, (b) y-log-scale gradient magnitude (top) and voxel value histogram (bottom) of slice. The high intensities of the metal artifacts can be seen on the far right end of the histogram.

$$I_c(x, y, z) = \frac{I(x, y, z) \left(1 + \frac{I(x, y, z)}{I_{MAX}}\right)}{1 + V_{i-1}(x, y, z, s)} \quad (9)$$

Since $I_c(x, y, z)$ generates intensities in the $[0; 1]$ interval, it must be scaled into the target 8bit range by scaling (and truncating) it with $(2^b - 1)$ with $b = 8$.

The described method is a local tone mapping operator, because the scale index i that is used for the actual mapping computation in Equation 9 is computed separately for every voxel and the respective filter has a limited local support. This has of course consequences for the data analysis, classification, and rendering, that we are going to discuss in Section 4.3.

4 RESULTS

The major achievement of volumetric high dynamic range (VHDR) windowing is a contrast enhancement in comparison to a standard linear windowing operation. In this section, we will discuss the use of the new windowing operators. Although we implemented also a slice-based variation of this operator, we will focus only on the volumetric version. While the computational costs per-voxel of the slice-based version are about 30% lower than of the volumetric variation ($O(n + n + n)$ vs. $O(n + n)$), the latter takes into account the full spatial information of the voxel surroundings.

In the course of this section, we will discuss Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) as two different imaging modalities. Unfortunately, no ultrasound or PET datasets were available at the time of the experiments. Hence, we cannot provide information on these modalities.

4.1 Computed Tomography

Figure 6a-c shows the same slice from a CT scan of a thorax, showing the lungs, the heart, the trachea, parts of the backbone, the ribs, and the sternum. In Figure 6a, the full active bit range of 12bits is mapped linearly to 8bits. However, the full data range is not used by the CT dataset; all voxels are located in the data value range of $[0; 2700]$ (or $[-1023; 1677]$ Hounsfield units (HU)). Furthermore, an upper interval of 1000 HU does not contain useful information. If we add a clip operation to remove unused voxel value ranges at the windowing operation, the upper end (in this case) of the new value range will show a significant increase in the histogram, suggesting that there is a material interface or another feature.

In contrast, VHDR windowing provides an adaptive correction of the voxel values (see Fig. 6b), depending on the overall intensity of the image and a correction that takes into account the local neighborhood (*dodging-and-burning*, see Section 3.2). At first glance, we can see that the overall image is brighter than the one with linear mapping. Pixel by pixel comparison actually revealed that the contrast between darker tissue areas and brighter bone areas increased. Figure 6c exposes this difference between linear and the VHDR windowing operator. While there is a general shift of intensity, higher differences are in areas with a high contrast (eg., between bone and dark tissue). This is of course what could be expected, since these areas are exactly those which are supposed to be enhanced by the method.

The second examined CT dataset is a scan of a backpack filled with a variety of objects. The dataset occupies the full 12bit range of $[0; 4095]$ (or $[-1023; 3071]$ Hounsfield units (HU)), whereas the metal artifacts caused by some of the objects ensure a rather dark intensity distribution for most objects (Fig. 6j). With VHDR windowing, the intensity contrast is significantly improved and considerably more structures – which were virtually invisible before – become visible (Fig. 6k). Figure 6l shows the voxels of one slice where the intensity changed most by our operator. Note that for this VHDR windowing, the key value a has been changed to 0.05 (default is $a = 0.18$).

4.2 Magnetic Resonance Imaging




Figure 2: Histograms of 3D CISS MRI scan of a head. (a) shows the standard histogram of the (stretched) linear (top) and VHDR windowing (bottom) result and (b) shows the respective gradient magnitude histograms. The windowing results are shown in Figure 6d-e.

Our second examined modality is MRI. In contrast to CT, MRI provides a whole zoo of different image acquisition protocols that generate image data focusing on different organs and tissue types. Furthermore, MRI generally provides more material (or tissue) differentiation than CT, thus making it an imaging modality with a good potential for contrast improvements. However, the different MRI image acquisition protocols generate data with limited 9-

10bits per voxel accuracy. Therefore, the data has a smaller dynamic range than datasets from CT.

In our first example, we examine a 3D CISS MRI imaging sequence, which is a T2-weighted protocol that emphasizes on the cerebro spinal fluid (CSF) filled cavities of the head. In contrast to our previous CT dataset, this MRI dataset has only 10 active bits and only the active value interval of $[0; 850]$. Hence, we only map 10bits to 8bits for all windowing methods (Fig. 6d-e).

Figure 2 shows the standard (voxel value) and gradient magnitude histograms. We stretched the histograms of the linear windowing (top histograms) to the full range, thus allowing an easier comparison. Overall, we can observe that the histograms of the VHDR windowing (bottom) are less affected by a small number of high intensity noise values (upper/right ends of linear histograms) and relevant regions have a smoother development of histogram values. In particular peaks are more spread out, thus enabling a better data analysis, segmentation, or classification of the data.

Also with the 3D CISS MRI sequence, the contrast differences are quite obvious between VHDR windowing and the linear windowing (Fig. 6f). Beside the general brightness shift between the dark background and the head, edges are again emphasized, since they represent high contrast differences. We can also observe the enhanced intensity of the soft brain tissue (gray and white matter) after VHDR windowing.

In our last example, we examine another MRI dataset. This dataset was acquired by T1-weighted flash sequence that mapped CSF-filled cavities to black (Fig. 6g-i). Similar to the previous 3D CISS MRI dataset, only 9bits are active in a voxel value interval of $[0; 350]$, hence only 9bits are mapped into 8bits by all windowing operators (Fig. 6g-i).

Again, we can observe a general brightness shift between the VHDR and linear windowing, and in addition a contrast enhancement on the edges and the enhanced soft brain tissue (Fig. 6h) intensity.




Figure 3: Image slice from MRI CISS dataset. (a) Original 8bit image slice, (b) VHDR enhanced 8bit image slice.

An interesting application of VHDR is contrast enhancement in low-fidelity datasets. Figure 3 demonstrates an example, where another 3D CISS MRI sequence dataset is only available in 8bit fidelity. Due to the low voxel fidelity, the appearance of the MRI image slice is very dark (Fig. 3a). After VHDR windowing (from 8bit to 8bit), the contrast is significantly enhanced (Fig. 3b), while the noise level is kept at a much lower level than with histogram equalization (see Fig. 5 for the MRI CISS example).

4.3 Discussion

In the past two sections, we presented the application of volumetric high dynamic range (VHDR) windowing to three different volumetric datasets of different modalities. We could observe a general contrast enhancement with VHDR windowing methods, compared to linear windowing. At the same time however, VHDR windowing can also increase the noise level, if high-intensity noise is present,

Table 2: Timings for the various windowing methods on five different volume datasets.

Dataset	Windowing Method	Time (s)
CT Head 512 x 512 x 324	Linear	2.94
	Zone Intensity	4.27
	VHDR Windowing	30.49
CT Thorax 512 x 512 x 168	Linear	0.17
	Zone Intensity	0.21
	VHDR Windowing	16.31
CT Angiography 512 x 512 x 194	Linear	0.57
	Zone Intensity	0.98
	VHDR Windowing	12.49
MRI CISS 512 x 512 x 54	Linear	0.06
	Zone Intensity	0.09
	VHDR Windowing	7.09
CT Backpack 512 x 512 x 373	Linear	1.06
	Zone Intensity	2.07
	VHDR Windowing	31.95

eg., with metal artifacts in CT data. For most practical aspects though, this increased artifact noise is negligible.

The increased contrast quality does not come for free; the full accuracy datasets need to be mapped using the linear or VHDR windowing operations in a pre-process. Not surprisingly, the required time depends on the size of the dataset and on the chosen windowing methods. In our experiments (see Table 2) with the various datasets discussed above, we measured timings between less than a second for linear windowing of the smallest datasets and up to 32 seconds for the full volumetric HDR windowing of the largest dataset. All measurements were performed on a PC, equipped with 2GB of main memory and an Intel P4 CPU running at 3GHz. We consider the increased computational costs for VHDR windowing as modest and easily affordable for the increased quality.

The VHDR windowing operators logically perform two processing passes to enhance the contrast in the datasets (in the actual implementation, both passes are integrated into one processing pass). While the first pass uses a global, logarithmic scaling of the datasets – thus transforming the histogram of the data –, the second pass applies a filter of the local voxel region around every voxel. This however, takes into account the local spatial neighborhood of the dataset and not only their voxel values, resulting in a spatially varying operation. As a consequence, previously available mapping methods based on the voxel value (eg., transfer functions) will affect a different set of voxels. Fortunately, the number of these voxels is small and they almost always belong to an edge area with a now increased contrast. Most changes have been introduced by the global zone intensity mapping, thus the transfer functions – based on the original histogram – can be adapted to the new voxel values, according to the transformation of the histogram. The changes introduced by the local *dodging-and-burning* are limited to areas with high contrast (edges) and only emphasize the material interfaces (see Fig. 4a-d). Therefore, the consequences for a previously known classification are manageable.

Figure 4c-f shows the visual differences based on a visual difference predictor (VDP). Figure 4e-f also indicates the differences between the slice-based and the volumetric HDR windowing operator.

We also compared the performance of VHDR windowing with other high dynamic range operators. In particular, we also tested Ashikhmin's tone reproduction operator [3] and Durand et al.'s operator that is based on bilateral filtering [9]. As it turns out, both operators are less suited for windowing than Reinhard et al.'s PTR operator [16]. Both operators in particular focus on high dynamic range images with high image luminances, which are not really




Figure 4: Visual differences between different windowing operators for CT Thorax (a,c,e) and 3D CISS MRI (b,d,f) datasets. The first row shows the relative intensity between zone intensity mapping and VHDR windowing (images are darkened for printing contrast). The other rows show a perceived intensity difference between zone intensity mapping and VHDR windowing (2nd row), and between slice-based HDR and VHDR windowing (3rd row) based on a Visual Difference Predictor [15]. (g) shows its color-coded scale of the detection probability.

present in medical image data. Figure 5c and d shows absolute intensity differences mapped to a hue colormap (blue no difference, red large differences). Most differences are located on the contrast edges of the dataset, with some differences also in the higher intensity areas of the CSF-filled ventricular system. This is not very surprising for Durand’s approach, since it is based on edge-preserving bilateral filtering. Our experience indicates that an operator which focuses more on the intermediate luminance values, like PTR, performs better for our application. Furthermore, Ashikhmin’s operator expects real luminance values for processing which cannot be substituted by the intensity values of the voxels, as it can be done with PTR. Finally, both operators are more costly to compute. Especially Durand’s bilateral filtering requires many costly *exp* operations, which considerably slow the algorithm down.




Figure 5: Other windowing operator applied to the MRI CISS dataset. (a) Original image slice, (b) histogram equalized image slice with increased background noise. (c) Absolute intensity difference between VHDR windowing (Fig. 6e) and after Ashikhmin’s tone reproduction operator and (d) Durand’s bilateral filtering.

4.4 Error Metrics

An important question is how we can measure if our VHDR windowing method is actually improving the image quality compared to the traditional linear windowing. The question which metric to use is quite complex. While most high dynamic range mapping methods can be evaluated by esthetics (“Is the mapped image more pleasing? Does it convey the information?”) or psycho-physical metrics, this is not straight-forward for volumetric medical imaging datasets. Linear windowing translates the intensities of the original data directly into the lower range. A simple error metric based on the normalized intensities would not detect contrast improvements achieved by our algorithms. Therefore, we decided to use an entropy-based metric, assuming that VHDR windowing will maintain more information of the original data than linear windowing. The entropy itself is calculated by $H = -\sum_i (p_i \cdot \log_2(p_i))$, with p_i as the probability of intensity i in a voxel.

Table 3 shows the results of the measured entropy of five datasets with our three different windowing methods, compared to the original volume dataset. For all three datasets, the entropy is monotonically increasing with the quality of the windowing method. Generally, the entropy of an image would be maximized in a completely (histogram) equalized dataset, but this would in particular enhance the low intensity background noise, since a large number of background voxels occupies only a small number of low intensity histogram buckets. Figure 5a and b shows a histogram equalized slice from the MRI CISS dataset with such an increased noise level, demonstrating why it is not suitable for the enhancement of scanned volume data. This example also clearly shows that the entropy-based metric is not a perfectly suitable metric. VHDR does some kind of equalization of the voxel values, since it redistributes the intensity values to enhance the intensity range representation.

As an alternative, we also adopted a contrast measuring metric that is used in multimedia texture analysis [17]. This contrast metric, described in Equation 10, represents the overall contrast in a

Table 3: Computed entropy for all windowing methods.

Dataset	Wind' Method	Entropy H	Contrast C
CT Head 512 x 512 x 324	Original 12bits	9.81	30331.20
	Linear	5.97	117.58
	Zone Intensity	6.33	148.94
	VHDR Wind'	6.52	266.77
CT Thorax 512 x 512 x 168	Original 12bits	9.76	2772.43
	Linear	5.76	10.92
	Zone Intensity	6.33	21.28
	VHDR Wind'	6.35	25.34
CT Angiography 512 x 512 x 194	Original 12bits	9.42	5126.89
	Linear	5.41	20.01
	Zone Intensity	6.29	30.98
	VHDR Wind'	6.35	43.69
MRI CISS 512 x 512 x 54	Original 10bits	8.01	921.27
	Linear	5.02	14.50
	Zone Intensity	6.92	123.91
	VHDR Wind'	7.00	208.55
CT Backpack 512 x 512 x 373	Original 12bits	4.55	8637.52
	Linear	2.43	33.44
	Zone Intensity	4.15	584.02
	VHDR Wind'	4.31	881.05

data volume. $g[i, j]$ represents the gray-level co-occurrence matrix, which enumerates the voxels with value i that have a neighbor with voxel value j , with H as the maximum voxel value of the respective histogram (eg., 255 for 8bit datasets), and N as the number of voxels. Note, that $g[i, j] \neq g[j, i]$, since different voxels have different neighborhoods. Table 3 provides also the results of the contrast metric for the different datasets. The significantly higher contrast values for the original dataset are due to the higher voxel depth of 12bits (10bits).

$$C = \sum_{i=0..H} \sum_{j=0..H} \frac{(i-j)^2 g[i, j]}{N} \quad (10)$$

Similar to the entropy metric, the calculated contrast C for the examined datasets is always significantly higher for the VHDR windowing methods, indicating an improved contrast.

5 CONCLUSION AND FUTURE WORK

In this paper, we presented a new windowing operator that is based on a recently introduced tone mapping method. This operator can map any data item size (eg., 12bit integer values) into any other data item size (eg., 8bit integer values) and optimizes its range representation to enhance contrast. The operator can also be used to map 8bit data into an 8bit data range, in order to enhance the intensity contrast.

We evaluated the performance of our operator with an entropy metric and with a contrast-based metric from video-encoding. Both metrics clearly showed a significant improvement with our operator. Furthermore, we employed a perception-driven visual difference predictor to expose areas of perceptually largest differences between different HDR operators.

Still, more research needs to be done to consider psycho-physical aspects of the human visual systems into an error metric. Another aspect of future work will look into the application of VHDR windowing for other data domains in medicine (eg., DTI datasets) or scientific visualization (eg., CFD). We believe that the methodology can also help here to map high dynamic range data into a more feasible range.

ACKNOWLEDGMENTS

This work has been supported by the Department for Neurosurgery of the University Hospital Tübingen and by the DFG project VIRTUE. Some datasets are courtesy of the Department of Radiology of the University Hospital Tübingen. The backpack dataset is courtesy of Kevin Kreeger and was downloaded from www.volvis.org. We thank Rafal Mantiuk of MPI Saarbrücken for providing us with his visual difference predicting software.

REFERENCES

- [1] A. Adams. The Camera. *The Ansel Adams Photography Series*, 1980. Little, Brown and Company.
- [2] A. Adams. The Print. *The Ansel Adams Photography series*, 1983. Little, Brown and Company.
- [3] M. Ashikhmin. A Tone Mapping Algorithm for High Contrast Images. In *Proc. of Eurographics Workshop on Rendering*, 2002.
- [4] Barco. MeDis Grayscale CRT Display Systems Product Brochure. Barco, 2003.
- [5] D. Bartz. Advanced Virtual Medicine: Techniques and Applications for Virtual Endoscopy. In *ACM SIGGRAPH Course 52*, 2002.
- [6] F. Blommaert and J. Martens. An Object-Oriented Model for Brightness Perception. *Spatial Vision*, 5(1):15–41, 1990.
- [7] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman. Spatially Nonuniform Scaling Functions for High Contrast Images. In *Proc. of Graphics Interface*, pages 245–253, 1993.
- [8] Kate Devlin. A Review of Tone Reproduction Techniques. Technical Report 117, Department of Computer Science, University of Bristol, November 2002.
- [9] F. Durand and J. Dorsey. Fast Bilateral Filtering for the Display of High Dynamic Range Image. In *Proc. of ACM SIGGRAPH*, pages 257–265, 2002.
- [10] R. Fattal, D. Lischinski, and M. Werman. Gradient Domain High Dynamic Range Compression. In *Proc. of ACM SIGGRAPH*, pages 249–256, 2002.
- [11] R. Gonzalez and R. Woods. *Digital Image Processing*. Addison-Wesley, Reading, 1992.
- [12] A. Kugler, T. Grunert, E. Becker, and W. Strasser. MEDStation: Bringing Together Medical Imaging and Diagnosis. In *Proc. of European Multimedia, Microprocessor Systems and Electronic Commerce*, 1998.
- [13] J. Mitchell. Real-Time 3D Scene Post-Processing. Game Developer Conference 2003: http://www.ati.com/developer/gdc/GDC2003_ScenePostprocessing.pdf, accessed 2004.
- [14] H. Petterson. Window (The Encyclopaedia of Medical Imaging website). <http://www.amershamhealth.com/medcyclopaedia/medical>, accessed 2005.
- [15] R. Mantiuk and K. Myszkowski and H. Seidel. Visible Difference Predictor for High Dynamic Range Images. In *Proc. of IEEE International Conference on Systems, Man and Cybernetics*, pages 2763–2769, 2004.
- [16] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic Tone Reproduction for Digital Images. In *Proc. of ACM SIGGRAPH*, pages 267–276, 2002.
- [17] R. Steinmetz. *Multimedia Technologie - Grundlagen, Komponenten und Systeme*. Springer-Verlag, Heidelberg, 3rd edition, 2000.
- [18] J. Tumblin, J. K. Hodgins, and B. Guenter. Two Methods for Display of High Contrast Images. *ACM Transactions on Graphics*, 1(18):56–94, January 1999.
- [19] J. Tumblin and G. Turk. LCIS: A Boundary Hierarchy for Detail-Preserving Contrast Reduction. In *Proc. of ACM SIGGRAPH*, pages 83–90, 1999.




Figure 6: Dataset overview: The left column shows the result from the linear windowing (a,d,g,j), the middle columns shows the results from the volumetric high dynamic range (VHDR) windowing (b,e,h,k), and the right column shows the relative voxel intensity difference between linear and VHDR windowing (c,f,i,l), where the occurring difference is mapped to the full hue range (blue no difference, red largest difference). The top row shows the CT thorax, the second row shows the MRI CISS dataset, the third row shows a dataset from an MRI flash dataset, and the bottom row shows results for the CT backpack dataset. The actual slice size of the flash MRI dataset is smaller, than the image size in Fig. 6g-i. This can be seen on the left and right side of the images.