Efficient Visualization of Volume Data Sets with Region of Interest and Wavelets

Sébastien Piccand[a], Rita Noumeir[a], and Eric Paquette[b]

[a]LIVIA, Electrical Engineering Department
École de Technologie Supérieure, Montréal, Canada
[b]LESIA, Software & IT Engineering Department
École de Technologie Supérieure, Montréal, Canada

ABSTRACT

The growing volume of medical images acquired with new imaging modalities poses big challenges to the radiologist's interpretation process. Innovative image visualization techniques can play a major role in enabling efficient and accurate information presentation and navigation, by combining computational efficiency with diagnostic resolution. Efficiency and resolution, two opposing requirements, can be accomplished by focusing on full resolution regions of interest while maintaining sufficient contextual information. In fact, structures of interest typically occupy a small percentage of the data, but their analysis requires context information like locations within a specific organ or adjacency to sensitive structures.

We propose a 3D visualization technique that is based on the multi-resolution property of the wavelet transform in order to display a full resolution region of interest while displaying a coarser context to achieve efficiency in rendering during the exploratory navigation phase. A full resolution context can also be rendered when needed for a specific view. In a preprocessing stage the data is decomposed with a three-dimensional wavelet transform. The interactive visualization process then uses the wavelet representation and a user-specified region to render a full resolution region of interest and a coarser context directly from the wavelet space through wavelet splatting, thus avoiding volume reconstruction. This efficient rendering approach is combined with lighting calculations, in the preprocessing stage. While greatly enhancing depth perception and objects shape, lighting does not add additional cost to the interactive visualization process, resulting in a good compromise between computational efficiency and image quality.

Keywords: Visualization, Rendering, Volume Rendering, Large Data Sets, Region of Interest, Wavelets, Splatting

1. INTRODUCTION

Patient care has been greatly improved with the recent deployment of digital medical images. Medical images are digitally archived, transferred on telecommunication networks, and visualized on computer screens. However, modern imaging acquisition modalities produce huge amounts of data. A magnetic resonance imaging (MRI) or a computerized tomography (CT) generate volume data composed of 100 to 300 images per exam, for a minimum of 50 Mbytes. Medical data sets are increasingly large. New medical imaging modalities, such as multi-slice spiral CT, provide substantial increase in image resolution. Very large data sets of more than 300 slices will be the standard in coming years. Functional MRI that offers simultaneous temporal and spatial information produces higher dimensional and thus larger data sets.

In general, the reading physician focuses on a region of interest (ROI) when interpreting medical images. The structure of interest, a tumor or an organ for example, often occupies less than 10% of all the data. Limiting the visualization to a region of interest reduces the amount of data to process, but sacrifices contextual information. Approaches that combine complete resolution ROI with a low resolution context promise to preserve

Further author information:
S.P.: E-mail: sebastien.piccand.1@ens.etsmtl.ca
R.N.: E-mail: rnoumeir@ele.etsmtl.ca
E.P.: E-mail: eric.paquette@etsmtl.ca

the diagnostic value of the data while reducing the computation time. We propose an approach that is based on a multi-resolution representation that uses the wavelet transform to integrate both efficiency and resolution for visualization of large volume data sets.

2. PREVIOUS WORK

Interpreting volume data is a complex task. It requires 3D visualization, which is achieved with different approaches. The simplest method is to display planar cuts through the volume, resulting in multi-planar images. Even though this visualization method is very accurate, the reading physician has to mentally reconstruct 3D objects of interest from 2D cuts. This process is variable, subjective, and may lead to incorrect decisions in diagnosis. A more complex approach extracts the object of interest out of the volume and only visualizes its surface.[1] In medical imaging, such extracted surfaces suffer from the lack of important contextual information. Alternatively, volume rendering[2,3] enables the observer to view all of the volume data simultaneously in one picture instead of perceiving either 2D sequentially displayed information or surfaces. This approach, which displays the volume in a translucent manner, presents all the 3D information. Volume rendering is based on considering the opacity of the voxels when integrating the volume along projection rays from the image plane. This approach is computationally intensive, especially with interactive manipulation of the volume. Thus, it can be time-consuming for a physician to wait for the system to render a new view of the volume data he is manipulating.

Because of the huge data sets to be handled, accelerated volume rendering methods are needed. Methods based on splatting[4] pre-compute the projection kernel with respect to the viewing angle. Object splatting,[5] has been used for skipping regions of no interest during voxel traversal. Hierarchical methods, such as the octree approach, have accelerated splatting by exploiting object-space coherence.[6]

The shear-warp factorization[7] technique is known as the most efficient software method to render large volume data sets. But the fastest techniques are now achieved with the use of graphics hardware.[8] The slices are sent to the card as textures and composed through the graphical processor unit.

Wavelets have provided very successful mathematical basis functions for representing and approximating signals because of their localization properties in both space and frequency domains. The wavelet decomposition has also proven to be very efficient for multi-resolution analysis.[9] The discrete wavelet transform decomposes the volume into a hierarchical subband system.[10] Wavelet decomposition of 3D volumetric images is achieved with 3D decomposition along the three axes.[11] The wavelet localization property has been greatly exploited in data compression.[12,13] Within the hierarchical subband system, every coefficient at a given scale is related to a set of coefficients at the next finer scale that correspond to the same spatial location in the image. This localization property allows for the efficient identification of all wavelet coefficients that compose the ROI. The use of a different resolution of the ROI has already been proposed but this ROI is fixed. If the ROI needs to be modified, the precomputation phase has to be made again. Recently, interactive modification of the ROI has been used,[14] but in order to be compatible with JPEG2000, block encoding is used instead of the localization property of wavelets. Wavelets have also been exploited in volume rendering,[15,16] to achieve progressive rendering of the complete data set. In this paper, we focus on visualizing a complete resolution ROI and a low resolution context, by processing all wavelet coefficients that compose the ROI and leaving out some wavelet coefficients that compose the context.

3. VISUALIZATION WITH REGION OF INTEREST

The method developed in this paper is summarized in Figure 1. For fast and accurate inspection of volumetric data presented as a 2D image, a ROI is used to discriminate between the essential and contextual region of the data. The volume is rendered at maximal quality within the ROI while an approximation is used when rendering the contextual region. Since the ROI usually represents a fraction of the whole volume, large gains in computation time are possible. In order to efficiently render the contextual region, the volume data is converted in a hierarchical multi-resolution representation in a precomputation phase. Given the ROI and the multi-resolution representation, a 2D image is computed with full detail in the ROI and only an approximation of the contextual region. This greatly reduces the time needed to render images in an interactive visualization of the volume data.

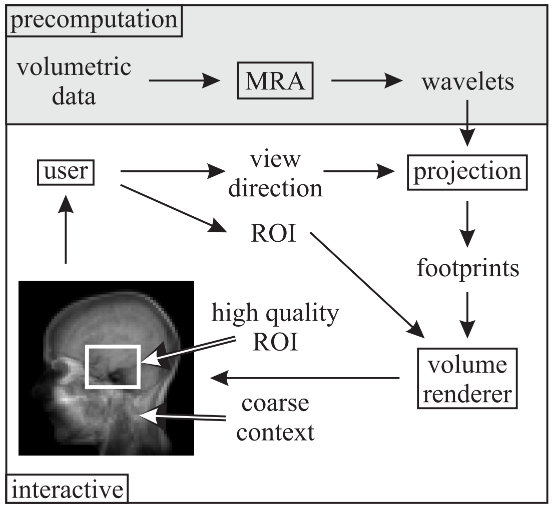

Figure 1. Summary of the method. In the precomputation phase, Multi-Resolution Analysis (MRA) is used to get a wavelet representation of the data. In the interactive phase, the input from the user is sent to both the wavelet projection computation and the volume renderer in order to compute the image with ROI.

3.1. Multi-resolution Representation

The volume rendering of the contextual region (outside of the ROI) will be an approximation at a coarser resolution. In order to compute this approximation efficiently, a multi-resolution representation of the volumetric data is used. The volumetric data, consisting of an opacity per volume element, is transformed in a wavelet representation in a precomputation phase. The wavelet representation reduces the computational expense of the interactive rendering. The use of wavelet does not increase the datasets size like a pyramidal decomposition would. We chose to use the wavelet of Haar because of its compact support in the spatial domain. It is therefore easier to define the coefficients that contribute to the region of interest.

3.2. Region of Interest

In an interactive visualization session, the user identifies a ROI that should be rendered at full resolution. The ROI is defined as an axis aligned box in the 3D space of the volume. To ease the perception of the ROI, a 2D ROI rectangle is computed from the 3D ROI and shown in the rendered image. To achieve the multi-resolution rendering of the ROI and contextual region, some wavelet functions will not be used. Every wavelet that projects inside the 2D ROI is considered, but wavelets that represent high frequency details outside of the 2D ROI are rejected. In the same visualization session, the user can refine or change the ROI to focus on other areas as shown in Figure 2.

We also propose the possibility to hide the data which lies behind or in front of the 3D ROI. This is particularly useful when some data hides the ROI. The Figure 3 shows a case where some bones hide the organ we want to observe.

3.3. Volume rendering

Volume rendering is computed with an X-Ray-like approach since it is fast and well suited for interactive rendering. It uses an orthogonal projection to compute the projection of a volume V. The projection equation calculated along the projection axis $\mathbf{n}$ (viewing direction) at position $\mathbf{x}$ (3D location corresponding to a pixel of the rendered image) is given by:

$$L(\mathbf{x}, \mathbf{n}) = \int_{S_E}^{S_L} V(\mathbf{x} + s\ \mathbf{n}) ds$$

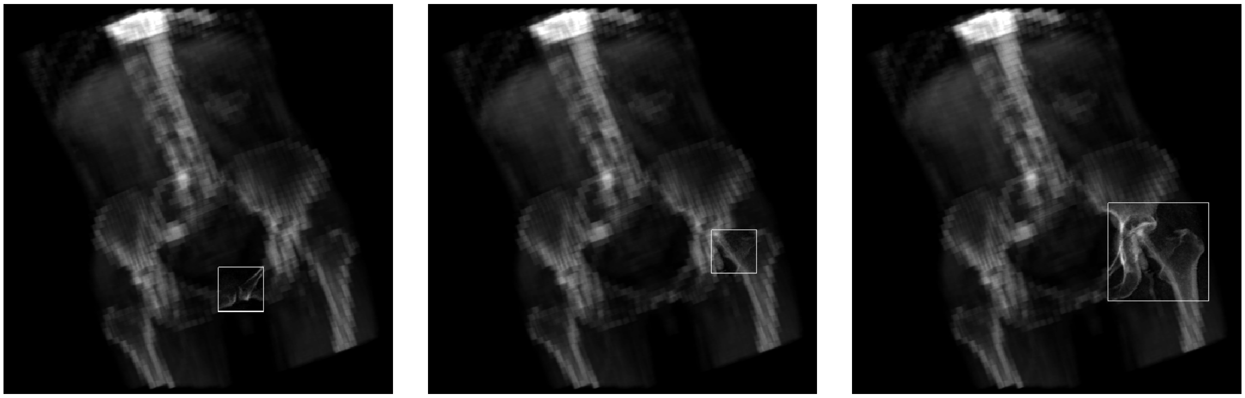

(a) (b) (c)

Figure 2. Interactive modification of the ROI.

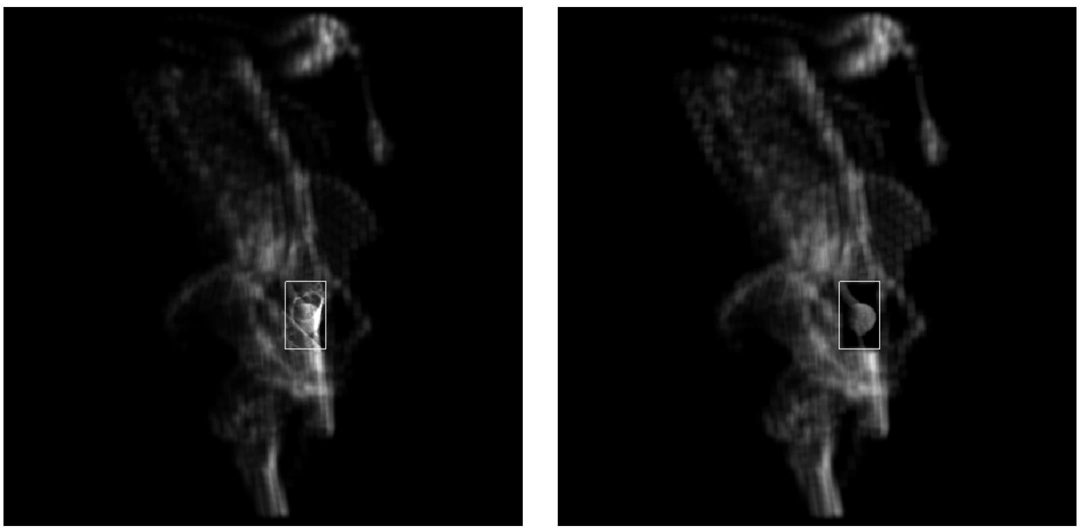

(a) (b)

Figure 3. Hiding objects in front of the ROI.

where S_E and S_L are the entering and leaving limits of the viewing ray respectively. The function $V(\mathbf{x} + s\ \mathbf{n})$ corresponds to sampling the appropriate pixel/slice of the volume. Since the volume V is decomposed by the wavelet transform along the 3 axis, it yields:

$$L(\mathbf{x}, \mathbf{n}) = \sum_{i \in LLL} c_i \int \Phi_i(\mathbf{x} + s\ \mathbf{n}) ds + \sum_{j \notin LLL} d_j \int \Psi_j(\mathbf{x} + s\ \mathbf{n}) ds \qquad (1)$$

where i indicates a position in the low frequency subband (LLL), Φ_i is the scaling function at that position and c_i are the low frequency wavelet coefficients ; j indicates a position in a high frequency subband, Ψ_j is the wavelet function at that position and d_j are the high frequency coefficients. This is a powerful representation of $L(\mathbf{x}, \mathbf{n})$ since it shows that the rendered image can be computed directly from the wavelet representation, at any appropriate level in the multiresolution representation, without reconstructing the volume. Thus, rendering directly from the wavelet representation is faster since it does not require the extra step needed to reconstruct the volume.

In order to render the volume from its wavelet representation, the wavelet splatting technique[17] is used. For every 3D wavelet function Φ and Ψ, a 2D projected *footprint* is computed prior to rendering. Each footprint corresponds to a fraction of the evaluation of the projection of Equation (1). This projection takes into account the direction of projection $\mathbf{n}$ and the particular function of the wavelet.

Since the wavelets are self-similar, the splatting integral is only computed for one function Φ and the seven types of wavelet functions Ψ.

Then when rendering the volume, each wavelet type is associated to the appropriate footprint which is multiplied by its coefficient as well as translated and resized to match the position and level of detail of the wavelet. The footprint is then added to the accumulation of the other footprints and the final image is the sum of all the wavelet footprints.

3.4. Shading

Volume ray-tracing[2] results in high quality images, but is computationally intensive. X-Ray rendering on the other hand is efficient, but removes the perception of depth and shape of the objects. Instead of using a plain X-Ray rendering approach, our method uses a compromise that precomputes shading in the volume. With shading precomputed, rendering is as efficient as the X-Ray approach, while presenting some important visual features of volume ray-tracing. Moreover using lightning attenuates the homogeneous parts of the volume which often do not give the most important information. Figure 4(a) and 4(b) show how precomputed shading eases the perception of depth and shape compared to the plain X-Ray rendering (Figure 4(c)).

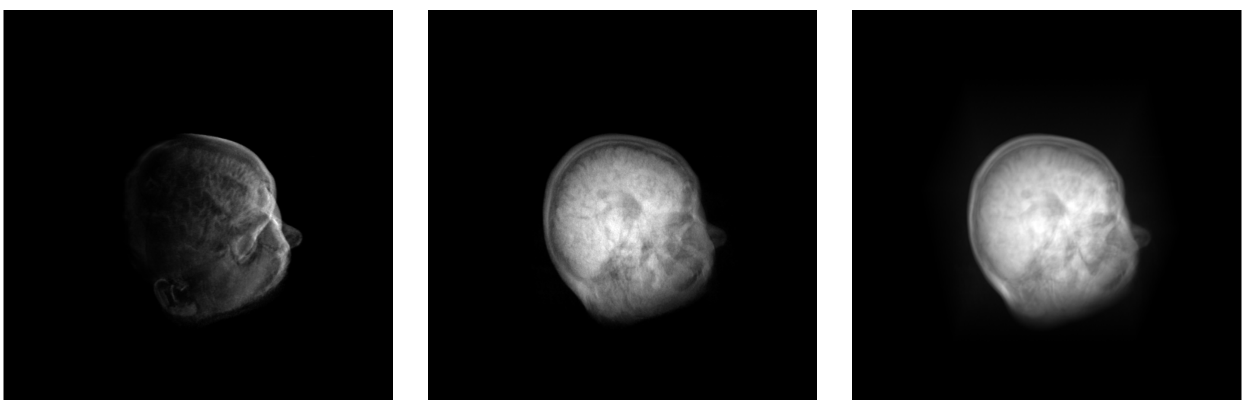

(a) Shadows and shading (b) Shading only (c) X-Ray

Figure 4. Increase of the perception of depth and shape with the use of precomputed lighting.

4. IMPLEMENTATION AND RESULTS

The proposed method is applied on MRI and CT data sets of size ranging from 256 to 512 slices. Using the Haar wavelet basis, the coefficients are obtained by decomposing the volume along the 3 axes into 8 subbands. The low frequency subband in the 3 directions is subsequently decomposed into 8 additional subbands. This is repeated once more. Figure 5 presents results for one data set. To illustrate the advantage of the proposed method, a ROI was chosen surrounding high frequency data. The projection was calculated by considering the lower frequency subband along the 3 axes in addition to all wavelet coefficients that compose the ROI. This projection is shown in Figure 5(b). For comparison, the full resolution is presented in 5(a) and rendering with only the low resolution coefficients is presented in 5(c). Figure 6 presents the same information for another dataset.

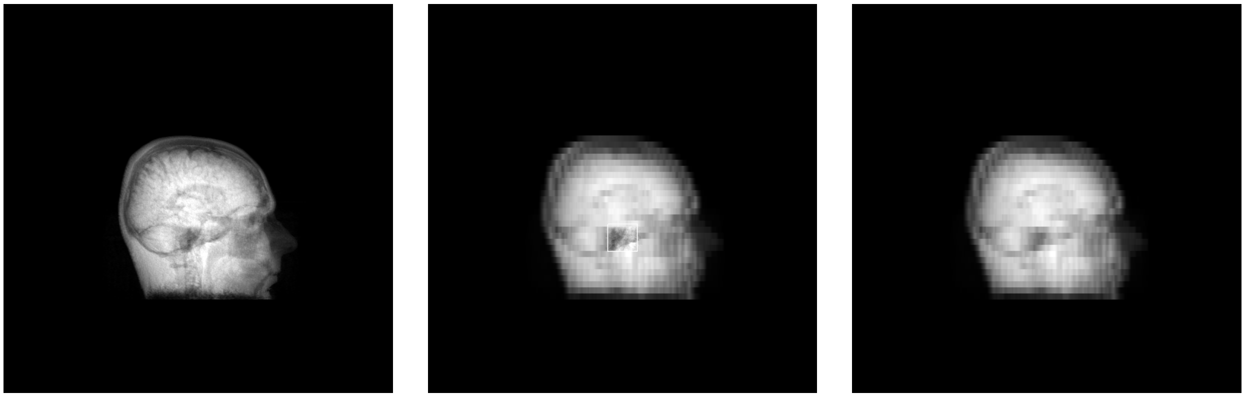

(a) Full (b) ROI (c) Low

Figure 5. 5(a) Rendering at full resolution. 5(b) Our approach renders the ROI (white square) at full resolution, while needing less computation time since the details of the context are left out. 5(c) For comparison, only using the low resolution coefficients is not sufficient for interpretation of the details in the ROI.

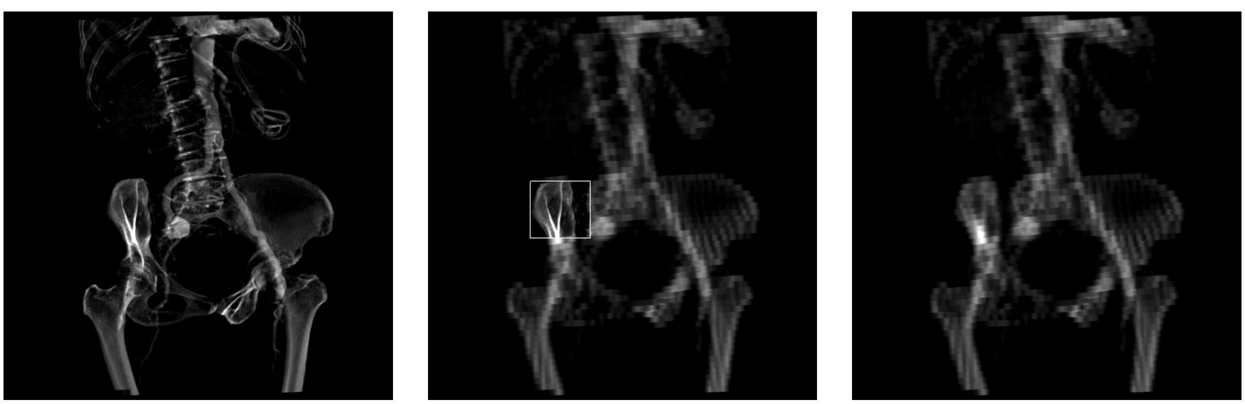

(a) Full (b) ROI (c) Low

Figure 6. Second dataset: 6(a) Rendering at full resolution. 6(b) Rendering with a ROI. 6(c) Rendering at low resolution.

From Equation (1), the projection to be displayed is obtained by summing the footprints which are used to precompute the integrals. By increasing the level of wavelet decomposition, the footprint dimension increases while the number of relevant wavelet coefficients decreases. By converting the data sets to wavelets, a gain in memory consumption can be accomplished since many coefficients are zeros. Table 1 presents statistics for the

Data set	Volume	ROI size	Number of coefficients			ROI speed-up
			Non-zero	context + ROI	low	
Fig. 5	256 × 256, 256 slices	36 × 36 × 36	45.34%	1.90%	0.14%	19
Fig. 6	272 × 472, 512 slices	38 × 66 × 72	4.6 %	0.33%	0.02%	19.7

Table 1. Statistics for the different data sets. The speed-up corresponds to the rendering time of the full resolution divided by the rendering time using ROI.

data sets presented in this paper.

Finally, the method is most useful in an interactive visualization context when a physician is interactively manipulating the volume data to interpret it. In this context, many images such as the ones of Figure 8 are computed with the reduced number of coefficients of the ROI approach (Figure 7(a)). By reducing the number of

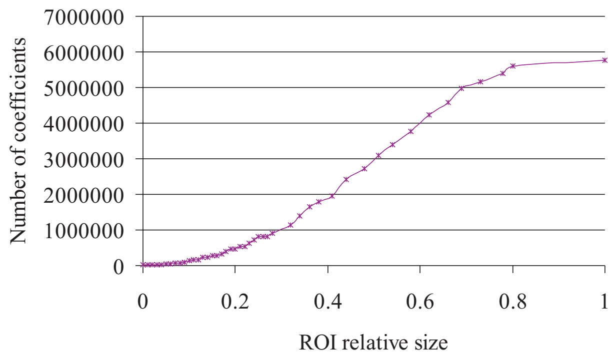

(a)

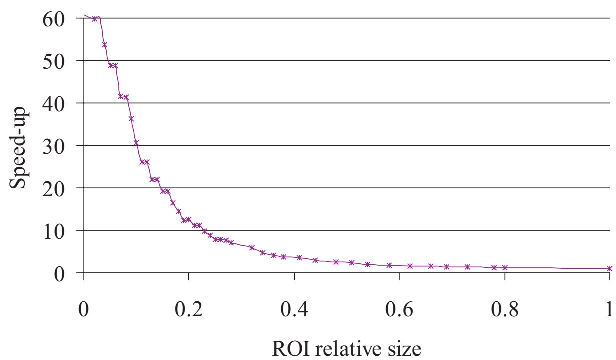

(b)

Figure 7. Impact of the ROI on the number of coefficients in 7(a) and on the speed-up of the rendering compared to full resolution rendering in 7(b). The ROI relative size is the size in pixels of the ROI relative to the size of the volume along one axis.

coefficients, the ROI also affects the rendering time (Figure 7(b)), so the visualization is much more interactive and the time needed to explore the data set is reduced.

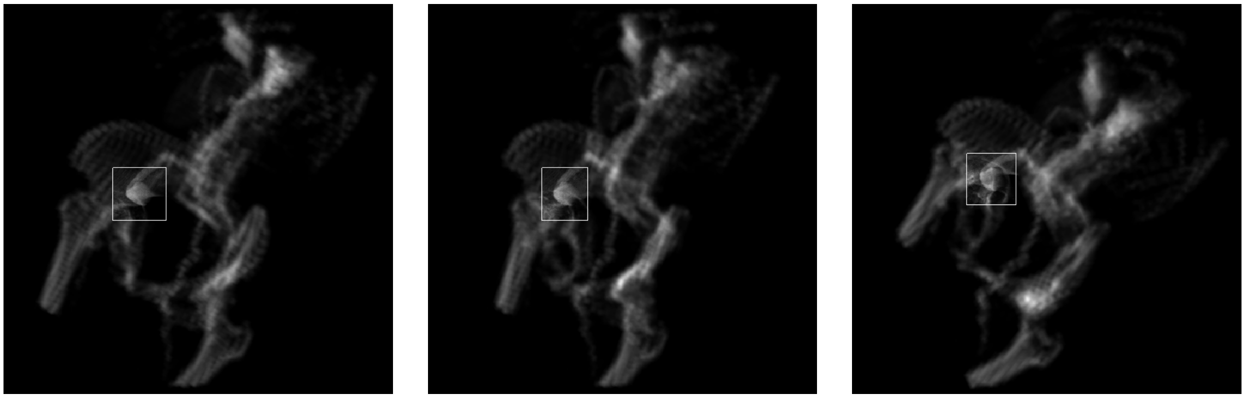

(a) (b) (c)

Figure 8. Images for a data set rendered with ROI for different viewing directions.

5. CONCLUSION

We propose a volume rendering method for interactive visualization of medical data. Our method is based on the 3D discrete wavelet transform. The breakthrough of our method is that it renders a full resolution region of interest and a lower resolution context. The ROI of our system is not fixed, allowing the physician to change its focus while interpreting the volume data. With the ROI simplification of the data and by rendering directly from the wavelet representation with wavelet splatting, it allows reductions in the computation time. The presented method is successfully used to render large volume data sets. It provides efficiency in computation time while ensuring the high resolution needed for medical image interpretation.

5.1. Future Work

One advantage of the presented method, is that it can be adapted to various types of wavelet basis other than Haar. With the Haar basis, the context rapidly becomes blocky. Other types of basis such as B-spline should reduce such disturbing artifacts. However using wavelets of higher order leads to larger footprints. The compromise between image quality and computing time should therefore be analyzed more closely. It would also be interesting to explore how much of the wavelet precomputation, footprint computation, and splatting could be computed on programmable graphics hardware. Significant gains in computation times are to be expected if all of the computation can be done on graphics hardware.

ACKNOWLEDGMENTS

This research was funded by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

1. W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *Computer Graphics (SIGGRAPH '87 Proceedings)*, **21**, pp. 163–169, July 1987.
2. J. T. Kajiya and B. P. V. Herzen, "Ray tracing volume densities," in *Computer Graphics (Proceedings of SIGGRAPH 84)*, **18**, pp. 165–174, July 1984.
3. M. Levoy, "Efficient ray tracing of volume data," *ACM Transactions On Graphics* **9**(3), pp. 245–261, 1990.

4. L. Westover, "Footprint evaluation for volume rendering," in *Computer Graphics SIGGRAPH Conference Proceedings*, **24**, pp. 367–376, August 1990.
5. R. K. Lee and I. Ihm, "On enhancing the speed of splatting using both object- and image-space coherence," *Graphical Models* **62**, pp. 263–282, July 2000.
6. H. Samet and A. Kochut, "Octree approximation and compression methods," in *Proceedings of the First IEEE International Symposium on 3D Data Processing Visualization and Transmission*, 2002.
7. P. Lacroute and M. Levoy, "Fast volume rendering using a shear-warp factorization of the viewing transformation," in *Proceedings of SIGGRAPH 94, Computer Graphics Proceedings, Annual Conference Series*, pp. 451–458, July 1994.
8. S. Guthe and W. Strasser, "Advanced techniques for high-quality multi-resolution volume rendering," *Computers & Graphics* **28**, pp. 51–58, february 2004.
9. S. G. Mallat, "A theory for multiresolution signal decomposition - the wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, pp. 674–693, 1989.
10. M. Vetterli and C. Herley, "Wavelets and filter banks: Theory and design," *IEEE Transactions on Signal Processing* **40**(9), pp. 2207–2232, 1992.
11. K. G. Nguyen and D. Saupe, "Rapid high quality compression of volume data for visualization," in *Computer Graphics Forum*, **20**(3), pp. 49–56, 2001.
12. A. Said and W. A. Pearlman, "A new fast and efficient image codec based on set partitioning in hierarchical trees," *IEEE Transactions on Circuits and Systems for Video Technology* **6**, pp. 243–250, 1996.
13. J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Transactions on Signal Processing* **41**(12), pp. 3445–3462, 1993.
14. K. Krishnan, M. W. Marcellin, A. Bilgin, and M. S. Nadar, "Compression/decompression strategies for large-volume medical imagery," *Medical Imaging 2004: PACS and Imaging Informatics* **5371**(1), pp. 152–159, SPIE, 2004.
15. M. H. Gross, L. Lippert, R. Dittrich, and S. Häring, "Two methods for wavelet-based volume rendering," *Computers & Graphics* **21**(2), pp. 237–252, 1997.
16. M. A. Westenberg and J. B. T. M. Roerdink, "An extension of fourier-wavelet volume rendering by view interpolation," *Journal of Mathematical Imaging and Vision* **14**(2), pp. 103–115, 2001.
17. D. Laur and P. Hanrahan, "Hierarchical splatting: A progressive refinement algorithm for volume rendering," in *Computer Graphics (Proceedings of SIGGRAPH 91)*, **25**, pp. 285–288, July 1991.