Exaggerated shading for volumetric data

Alark Joshi*
University of Maryland Baltimore County

Penny Rheingans†
University of Maryland Baltimore County

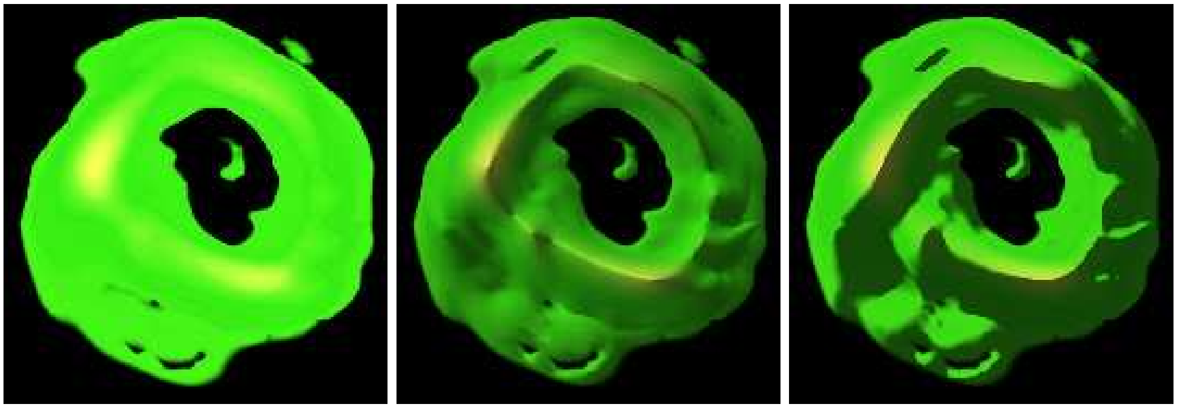

Figure 1: Volume rendered images of the cloud water attribute of a hurricane. The left image is a volume rendered image which shows the structure of the eye of the hurricane and surrounding eyewall. The central image shows a volume rendered image with lighting. The image shows a clear boundary of the eyewall and more structural details can be seen. The rightmost image shows a volume rendered image with exaggerated shading. The image depicts an opening in the eyewall towards the left side which is not at all visible in either of the other images. Such information can be critical to predicting the dissipation of a hurricane.

ABSTRACT

Illustrations convey overall shape as well as surface detail using certain lighting and shading principles. We investigate the use of such an illustration-inspired lighting model to accentuate features automatically by dynamically modulating the light position.

We apply this shading technique to accentuate detail in volumetric data. As the technique is primarily a gradient-based technique, we discuss and compare gradient computation techniques and their effectiveness. The technique highlights details at various levels in the volume making use of the multiple levels of blurred gradients computed. The results demonstrate that surface detail is accentuated regardless of the surface orientation and the size of features. We have applied our techniques to scientific and medical data to accentuate surfaces features in the application domains.

1 INTRODUCTION

Scientific illustrations have the ability to convey information based on the intent of the illustration. Illustrators have been using lighting and shading principles to convey overall shape information as well as finer details regarding the surface of the subject.

Automatically generating illustrative images from scientific data is a challenging task. Applying illustration-inspired principles to visualizing scientific data can produce expressive and insightful images. Generating effective visualizations of large datasets can be aided by the use of such techniques that provide context as well as are able to provide enhanced surface detail to the viewer.

Recent work in the field of illustrative computer graphics has dealt with generating feature-enhanced images by applying techniques from the field of mapmaking and scientific illustrations [Rusinkiewicz *et al.*, 2006]. They observed that artists and illustrators accentuate the lighting and shading to highlight regions of interest and draw the attention of a viewer to a certain region. Their technique does not require any explicit feature detection and automatically enhances features in the data.

Some of the illustration principles for shaded relief, as discussed by cartographers, that are applicable are [Rusinkiewicz *et al.*, 2006]:

- There should not be any specular highlights or shadows in the illustration.
- The light position should seem to be towards the upper left corner of the subject being illustrated.
- The light direction is locally adjusted to best emphasize local features.
- The height of features should be exaggerated and the transitions between ridges and valleys should be accentuated.
- In order to accentuate local as well as global features, the blending between various smoothed levels is used.

In the field of volume visualization, such techniques would prove invaluable as the datasets are large as well as highly detailed. For example, in a medical CT scanned dataset not only is the size of a dataset huge, but the detail of each anatomical part is crucial. Visualizing such data for exploratory as well as diagnostic purposes is quite challenging. The use of standard techniques seems to produce either a soft, diffused visualization that conveys overall shape information without detail, or a crisp detailed zoomed-in visualization where the context is lost.

The *exaggerated shading* technique with some modifications can be applied to generate expressive images of volumetric data. For volumes, we use the gradient of a voxel as the normal of the surface. As the effectiveness of the technique is based on the quality of

*e-mail: alark1@cs.umbc.edu
†e-mail:rheingan@cs.umbc.edu

the computed gradients, we investigate the use of the linear regression based gradient computation technique that has been proven to produce better gradients than the central differences method.

The exaggerated shading techniques are successful at accentuating regions due to their dynamically adjusted lighting direction. This is an excellent alternative to the geometry dependent lighting approach [Lee *et al.*, 2005] where a light source is placed in regions of interest to emphasize those regions to draw the viewer's attention. This technique automatically exaggerates the surface detail without requiring a manual/automatic placement of light sources around each feature. This technique helps in generating expressive images in an illustrative manner without any user input and provides the user with context as well as surface detail in the visualization. Expressive images generated in this manner will facilitate a better exploration of the dataset and lead to increased insight.

2 Related Work

Illustrative visualization is a field of wide interest and there have been many different approaches to generating expressive and illustrative images. Rheingans and Ebert [2001] have used gradient-based techniques to identify and accentuate features in a volume. Kindlmann et al. [2003] introduced the ability of producing insightful images with the use of curvature based techniques to enhance volumetric data automatically.

In volumetric datasets, the gradient serves as the normal for the surface and is used for lighting, shading as well as enhancement purposes. Due to its simplicity and easy of use, the central differences method is widely used in volume visualization to compute gradients of a volume. Linear regression based gradient computation techniques [Neumann *et al.*, 2000] though computationally expensive are better at computing the gradient for volumetric data.

Volume Rendering has been performed primarily using one of raycasting [Drebin *et al.*, 1988], shear-warp [Lacroute and Levoy, 1994], texture-mapping based techniques [Cabral *et al.*, 1994] or splatting [Westover, 1990]. There have been discussions regarding the trade-off between image quality and interactivity, where researchers have preferred texture-mapping based volume rendering for its speed and raycasting for the increased image quality. Our aim is to generate expressive images and so we chose the raycasting approach to generate higher quality expressive images.

3 Approach

The original exaggerated shading paper [Rusinkiewicz *et al.*, 2006] enhances ridges and valleys by using an illustrative lighting model. It relies on preserving and accentuating the edges as they are essential to a complete and correct understanding of the underlying data.

The technique blurs the normals at various levels and in the final calculation considers the contribution of each level. Their technique uses Gaussian blurring to produce multiple levels of blurred normals. In this process of blurring though, the blurring is performed without regard for boundaries between these rides and valleys. Bilateral filtering has been known to preserve edges and boundaries in subsequent blurred representations. We used a three dimensional version of the biltateral filtering algorithm discussed by [Tomasi and Manduchi, 1998].

Some other limitations of the original paper [2006] are that they used a texture-mapping based approach for their volume rendering example. In case of texture mapping based volume rendering, the volume rendering integral is evaluated by inserting slices orthogonal to the view vector and subsequent blending. In the raycasting approach, the volume rendering integral is solved along the ray as the ray traverses the volume. The difference between the two is that during ray casting the sample spacing along the ray is constant whereas in texture-mapping based volume rendering the slices lead to an increase in the sampling distance based on the view vector. Raycasting has always been known for better image quality and led to our decision of using it for our adaptation of the exaggerated shading technique.

As the technique relies heavily on the quality of the computed normals, use of the central differences method for computing the gradients is limited due to the small neighborhood that it searches during its computation. The 4D linear regression techique was introduced by Neumann et al [2000] and has been proven to more accurately compute gradients of volumetric data. We demonstrate and compare the results of these gradient computation techniques in our exaggerated shading framework.

3.1 The Shading Algorithm

As per the technique introduced in the original paper [Rusinkiewicz *et al.*, 2006] paper, the shading algorithm consists of the local lighting model, multiscale shading and modified local lighting to accentuate local features.

The local lighting model is a modified cosine shading model with an ambient term in order to prevent any self shadowing.

$$\frac{1}{2} + \frac{1}{2}\hat{n} \cdot \hat{l} \tag{1}$$

The light source is placed in the upper left corner of the image and $\hat{n}$ denotes the gradient direction of the voxel under consideration.

Decaudin [1996] proposed a cartoon style lighting model where the results are weighted to be closer to pure white and black. The paper discusses the fact that soft toon shading is much more effective at enhancing the contrast between ridges and valleys. In the original paper exaggerated shading paper they used a "soft toon shader" that was based on these principles and is given by,

$$\frac{1}{2} + \frac{1}{2} clamp_{[-1..1]} a(\hat{n} \cdot \hat{l}) \tag{2}$$

where a is a user-selected "exaggeration" parameter. Figure 2 shows results of cosine lighting on the left image followed by three toon-shaded images of the sphere with various settings for parameter a.

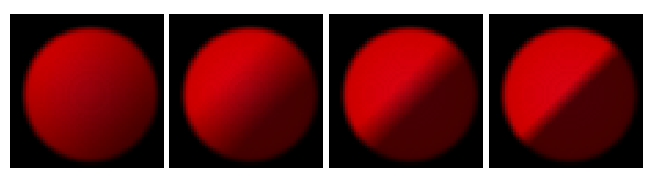

Figure 2: The leftmost image depicts a volume rendered image of a cosine lit model of a sphere volume. The next three images depict a toon-shaded image of the sphere with increasing values of the exaggeration parameter a =2, 4 and 8 respectively.

Local lighting is the third component of exaggerated shading. Instead of relying on a single light source to highlight all the features or placing multiple light sources to highlight specific features, in this technique the light vector is modified based on the gradient of the voxel to enhance local features.

The soft toon shading is computed using a modified version where

$$c_i = clamp_{[-1..1]} a(\hat{n}_i \cdot \hat{l}_{i+1}), \tag{3}$$

where

$$l_{i+1} = l_{global} - \hat{n}_{i+1}(\hat{n}_{i+1} \cdot l_{global}). \quad (4)$$

In the above equations, c_i is the color of the voxel after lighting and shading, a is the exaggerated shading parameter, $\hat{n}_i$ is the gradient of the voxel at the current scale, $\hat{n}_{i+1}$ is the gradient obtained by smoothing the gradient at level i. l_{global} is the light vector and l_{i+1} is the modified light vector taking the smoothed normals into account. We discuss the process of smoothing normals and generating multiple levels of smoothed normals in the next section.

3.2 Multiscale shading

Large datasets have information at various levels of detail. For example, a cartographic dataset will have structural information about an entire region, but at the same time the detail in a valley may be of interest to a user and within that valley, there maybe more relevant information regarding the surface of the valley. Such multiscale detail cannot be easily captured in a single visualization using standard visualization techniques.

The paper proposes multiscale shading by using the multiple levels of blurred normals. For multiscale shading, the above lighting calculations are performed with different levels of smoothed gradients for lighting computation. The idea is to actually compute the light based on different smoothed levels of the surface, but in practice smoothed versions of the gradient are used.

The smoothing of normals can be performed using a simple gaussian smoothing algorithm around each voxel. The bilateral filtering technique [Tomasi and Manduchi, 1998] introduced for images has been known to preserve the edges while filtering out detail. It can be formulated as [Fleishman *et al.*, 2003]:

$$I'(u) = \frac{\sum_{p \in V(u)} e^{-\frac{\|u-p\|^2}{2\sigma_c^2}} e^{-\frac{|I(u)-I(p)|}{2\sigma_s^2}} I(p)}{\sum_{p \in V(u)} e^{-\frac{\|u-p\|^2}{2\sigma_c^2}} e^{-\frac{|I(u)-I(p)|}{2\sigma_s^2}}}$$

In the equation for bilateral filtering, the first exponential term in the numerator filters out noise from the data, the second exponential term in the numerator ensure preservation of features. The second exponential term in the denominator is for normalization.

We have extended this technique to three-dimensions for blurring the gradients to generate multiple levels of gradients for the exaggerated shading technique. Figure 3 shows different levels of gradients obtained using gaussian blurring and bilateral filtering. As can be seen from the images, the bilateral filtering clearly smooths out more of the noise in the higher levels but is able to preserve the features surrounding the eyes and nose even at higher filtered levels.

As can be seen in figure 4, the images depict a volume rendered image of the engine dataset obtained by the use of multi-scale shading. For the left image, the multiple levels of gradients were generated using gaussian smoothing whereas for the right image, the multiple levels were generated using bilateral filtering. A marked difference can be seen in the way the lower left tube of the engine is rendered. More detail is seen in the left image that was generated using gaussian blurring. The right image shows a more smoothed version in that particular region. Features such as the indentation going across the engine is clearly seen in the right image as compared to the left image.

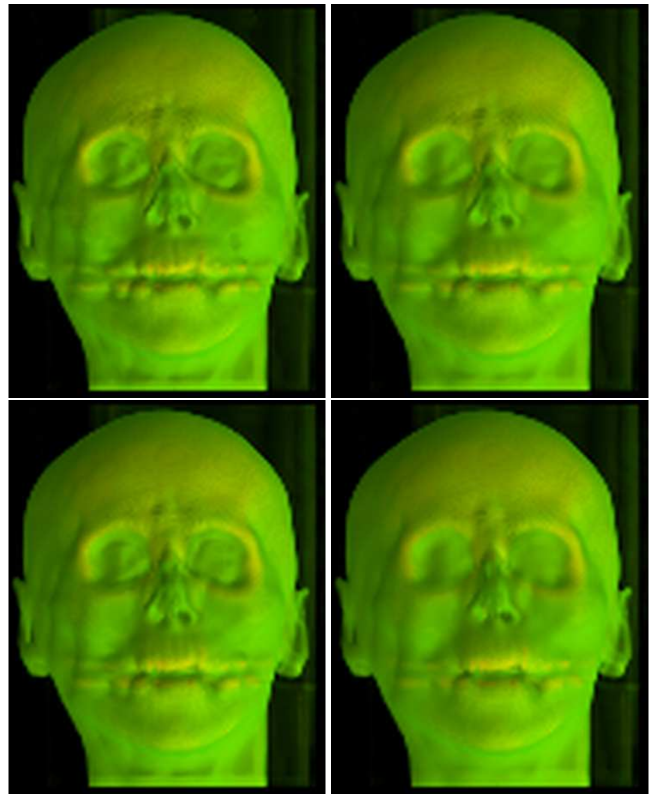

Figure 3: The top row contains volume rendered images of gaussian blurred second and third level gradients. The bottom row contains filtered gradients for second and third level obtained by bilateral filtering. The noise in the bottom row is better filtered than the top row. The features such as the eye sockets and the nostrils are better preserved in the bottom row as compared to the top row.

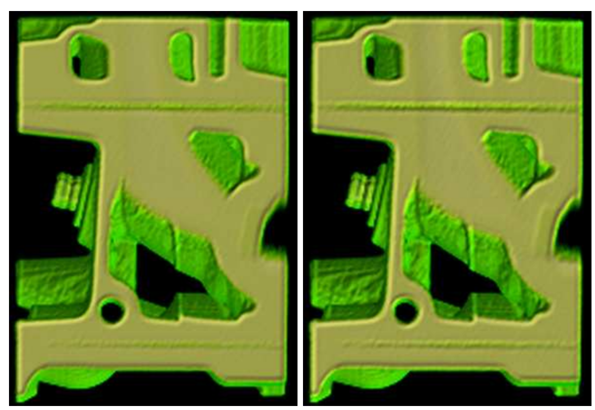

Figure 4: These images depict a volume rendered image of the engine dataset with exaggerated shading. The left image was generated by using gaussian blurring of the gradient at different levels and the right image was depicted by using bilateral filtering for blurring the gradients.The right image clearly depicts surface features better such as the curving indentation on the engine's surface and accentuates the boundaries much better than the left image.

3.3 Gradient computation techniques

The exaggerated shading technique is based on computing the accuracy of the gradients, we investigate the use of linear regression based gradient computation in comparison with the standard central differences method. The central differences method computes gradients by comparing voxel values along each dimension. The equation for computing the gradient using central differences is as follows:

$$\nabla_f(x_i,y_i,z_i) = \begin{bmatrix} \frac{1}{2}f(x_{i-1},y_i,z_i) - \frac{1}{2}f(x_{i+1},y_i,z_i) \\ \frac{1}{2}f(x_i,y_{i-1},z_i) - \frac{1}{2}f(x_i,y_{i+1},z_i) \\ \frac{1}{2}f(x_i,y_i,z_{i-1}) - \frac{1}{2}f(x_i,y_i,z_{i+1}) \end{bmatrix}$$

where $\nabla_f(x_i,y_i,z_i)$ is the gradient at a location P_i and $f(x_i,y_i,z_i)$ is the value at that location.

Neumann et al.[Neumann *et al.*, 2000] introduced a 4D linear regression based method for gradient estimation which approximates the density function in its local neighborhood with a 3D regression hyperplane. Their techniques produces smoother, high-quality gradients that can be used for lighting, shading and other gradient-based techniques for enhancing volume visualization.

A spherical function is used to weight the contributions of neighboring voxels in the computations of the plane equation coefficients. The plane equations are:

$$A = w_A \sum_{k=0}^{26} w_k v_k x_k \qquad B = w_B \sum_{k=0}^{26} w_k v_k y_k$$

$$C = w_C \sum_{k=0}^{26} w_k v_k z_k \qquad D = w_D \sum_{k=0}^{26} w_k v_k$$

where

$$w_A = \frac{1}{\sum_{k=0}^{26} w_k x_k^2}, w_B = \frac{1}{\sum_{k=0}^{26} w_k y_k^2}$$

$$w_C = \frac{1}{\sum_{k=0}^{26} w_k z_k^2}, w_D = \frac{1}{\sum_{k=0}^{26} w_k}$$

and the weighting function w_k is inversely proportional to the distance d_k of the neighboring voxel from the current voxel. The voxel gradient is then given by

$$\nabla_f = \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

We have previously experimented with these gradient computation techniques and found the 4D linear regression techniques to be smoother and of higher quality [Rheingans and Ebert, 2001].

Figure 5 shows volume rendering with exaggerated shading applied to an engine dataset. The left image shows the result obtained by using central differences, whereas the right image was obtained with gradient computed using 4D linear regression. The left image shows more of the surface detail of the flat surface of the engine compared to the right image. The right image gives more information regarding the depth of features, especially the big hole situated in the upper right region of the engine.

4 HARDWARE ACCELERATED EXAGGERATED SHADING

With the ever-increasing capabilities of newer graphics cards, performing hardware accelerated raycasting has become possible. We implemented the exaggerated shading technique in the fragment shader to analyze its impact on the overall interactivity of the application. We performed some experiments with datasets of various sizes and found a consistent 3 to 4 frames per second impact on the overall interactivity regardless of the size of the dataset. The results are reported in Table 1. The performance of the hardware accelerated raycasting decreases by 25-40% as regards the frames per second. As regards the percentage increase in the number of instructions in a shader that a graphics card has to execute, the shader's length increases by about 15%.

We believe that such results make the technique usable for real-world medical and scientific applications with large datasets.

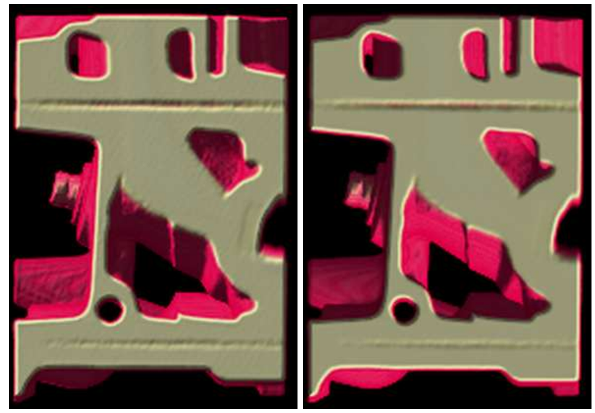

Figure 5: These images depict volume rendered images of the engine dataset with exaggerated shading. The left image was generated using central differences method for gradient computation. The right image was generated using the 4D linear regression method for gradient computation. The gradients generated by the linear regression method are much smoother and highlight important features such as the hole in the engine towards the bottom.

Name	*Dimensions*	*Raycasting*	*Lit*	*Exaggerated*
Sphere	64x64x64	14	12	8
Engine	128x128x64	13	12	7
UNC Brain	128x128x72	13	12	7
Tooth	128x128x256	13	12	8
CT Head	256x256x112	12	10	7
Abdomen	128x256x256	12	11	7
Human Head	256x256x256	12	11	5

Table 1: Impact of exaggerated shading on hardware accelerated raycasting

5 RESULTS

We applied the exaggerated shading technique to various volumes from different domains. In the medical domain, the UNC brain dataset was used as a representative dataset, since it has complex detail of the brain that needs to be visualized. Figure 6 shows a volume rendered image without and with exaggerated shading. Our discussions with doctors have confirmed that such visualizations of the gyri (folds) and the sulci (valleys) of the brain can help in diagnosing the damage caused due to a brain stroke. They said that it could also help in determining the extent of a brain tumor that needs to be surgical removed.

On applying our techniques to CT scanned abdomen data, we found the exaggerated shading technique to effectively accentuate the surface features in the data. Figure 7 shows an image generated using exaggerated shading where the bowl-shaped indentation in the hip bone (illiac fossa) as well as detail of the stent on the abdominal aorta can be clearly seen. Figure 8 shows a comparison of the use of bilateral filtered and Gaussian smoothed gradients. The left image shows bilateral filtered gradients that seem to highlight the left illiac fossa better than the right image which is obtained using Gaussian smoothed gradients.

We applied our technique to the domain of hurricane visualization and visualized the cloud water attribute from hurricane Katrina. Figure 1 shows our results on applying exaggerated shading. Only the rightmost image with the exaggerated shading shows a break in the eyewall. Such a discontinuity can be helpful to scientists to predict the dissipation of a hurricane.

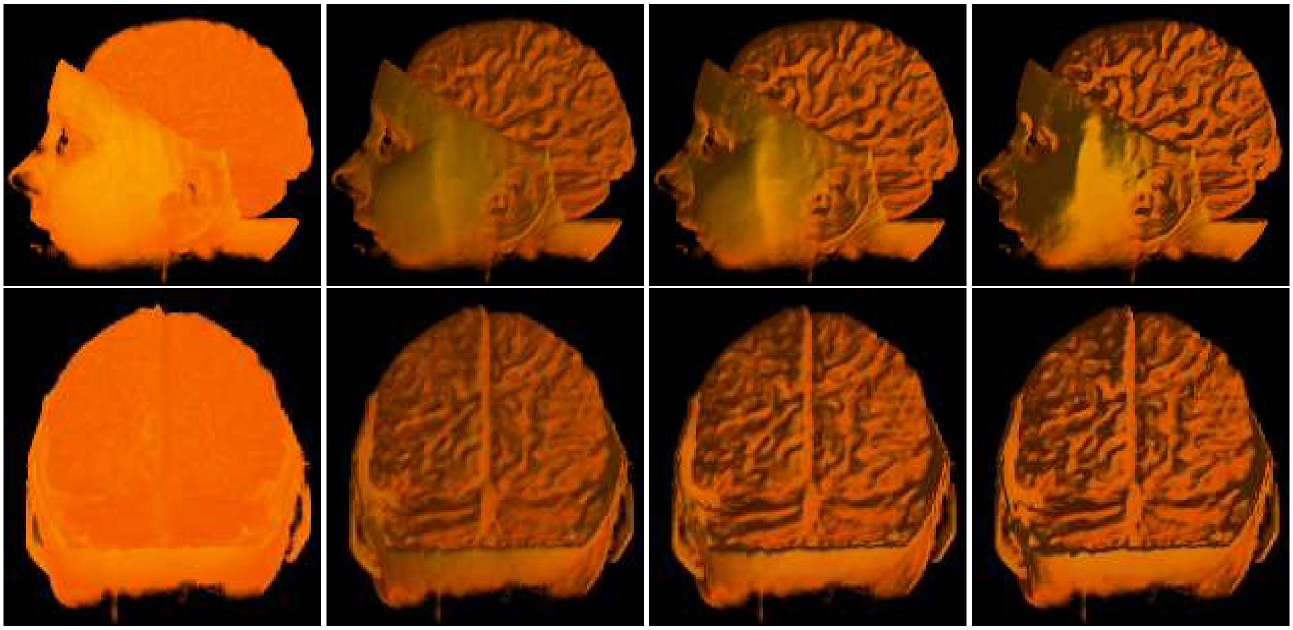

Figure 6: Volume rendered images of the UNC Brain dataset. The top row shows volume rendered images from the side and the bottom row shows images rendered from the back. The leftmost images are raycasted without lighting, the second column contains volume rendered images with lighting and the third and fourth column shows images with exaggerated shading parameter a=2 and a=8. The exaggerated shaded images accentuates the ridges and valleys in the brain. Clearly visualizing the gyri (folds) and the sulci (valleys) of the brain can help in diagnosing the damage caused due to a brain stroke.

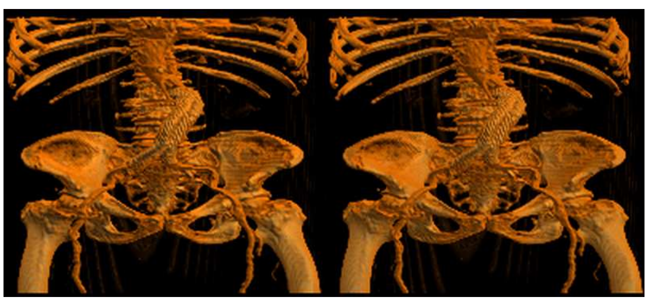

Figure 8: These images show a comparison of the bilateral filtering and the Gaussian smoothing techniques being used in the exaggerated shading technique with parameter $a = 4$. The left image better depicts the illiac fossa in the image with bilateral filtering than the image with gaussian smoothing in the bottom right.

6 Contributions

Our paper discusses the use of illustration-inspired lighting and shading for volumetric datasets. The main contributions of this paper are the discussion regarding the use of raycasting instead of the texture-mapping approach, the comparison between different gradient computation techniques and the benefits of using bilateral filtering instead of Gaussian filtering for the generation of multiple levels of the gradients.

7 Conclusion

We have applied the exaggerated shading technique to enhance volumetric features. We have demonstrated the applicability of the techniques for the medical visualization domain as well as the hurricane domain. The exaggerated shading technique highlights features of interest without the need for placing multiple local light sources and shows more information than using standard volume rendering techniques for visual investigation of a dataset. The gradient computation can be performed using 4D linear regression instead of central differences to generate better gradients. The three-dimensional versioan of bilateral filtering generates smoother representations of the gradients at multiple scales while preserving the boundaries and edges in the smoothing process.

References

[Cabral *et al.*, 1994] Brian Cabral, Nancy Cam, and Jim Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proceedings of the 1994 symposium on Volume visualization*, pages 91–98. ACM Press, 1994.

[Decaudin, 1996] Philippe Decaudin. Cartoon looking rendering of 3D scenes. Research Report 2919, INRIA, June 1996.

[Drebin *et al.*, 1988] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 65–74. ACM Press, 1988.

[Fleishman *et al.*, 2003] Shachar Fleishman, Iddo Drori, and Daniel Cohen-Or. Bilateral mesh denoising. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 950–953, New York, NY, USA, 2003. ACM Press.

[Kindlmann *et al.*, 2003] Gordon Kindlmann, Ross Whitaker, Tolga Tasdizen, and Torsten Moller. Curvature-based transfer functions for direct volume rendering: Methods and applications. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 67, Washington, DC, USA, 2003. IEEE Computer Society.

[Lacroute and Levoy, 1994] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 451–458. ACM Press, 1994.

[Lee *et al.*, 2005] C. H. Lee, X. Hao, and A. Varshney. Geometry-dependent lighting. *IEEE Transactions on Visualization and Computer*

Figure 7: These images show volume rendered images of a CT scan of the abdomen and pelvis. The patient has an a stent on the abdominal aorta which can be seen in the images. The left two images contain the volume rendered images of the dataset without and with lighting. The stent can be clearly seen and is higlighted. The rightmost image show the exaggerated shading techniques with parameter $a = 4$. The bowl-shaped indentation in the hip (illiac fossa) is better seen in the exaggerated shading image. The dataset was obtained courtesy of volvis.org.

Graphics, 12(2):197–207, 2005.

[Neumann *et al.*, 2000] Laszlo Neumann, Balázs Csebfalvi, Andreas König, and Meister Eduard Gröller. Gradient estimation in volume data using 4d linear regression. In *Proceedings of Eurographics*, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2000. human contact: technical-report@cg.tuwien.ac.at.

[Rheingans and Ebert, 2001] Penny Rheingans and David Ebert. Volume illustration: Non-photorealistic rendering of volume models. In *IEEE Transactions on Visualization and Computer Graphics*, volume 7(3), pages 253–264, 2001.

[Rusinkiewicz *et al.*, 2006] Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3), July 2006.

[Tomasi and Manduchi, 1998] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, page 839, Washington, DC, USA, 1998. IEEE Computer Society.

[Westover, 1990] Lee Westover. Footprint evaluation for volume rendering. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, pages 367–376. ACM Press, 1990.