Visualization of Noisy and Biased Volume Data Using First and Second Order Derivative Techniques

M.P. Persoon[1,2,3], I.W.O. Serlie[2], F.H. Post[1], R. Truyen[3], F.M. Vos[2,4]

[1]Computer Graphics and CAD/CAM Group, Delft University of Technology

Abstract

The quality of volume visualization depends strongly on the quality of the underlying data. In virtual colonoscopy, CT data should be acquired at a low radiation dose that results in a low signal-to-noise ratio. Alternatively, MRI data is acquired without ionizing radiation, but suffers from noise and bias (global signal fluctuations). Current volume visualization techniques often do not produce good results with noisy or biased data.

This paper describes methods for volume visualization that deal with these imperfections. The techniques are based on specially adapted edge detectors using first and second order derivative filters. The filtering is integrated into the the visualization process.

The first order derivative method results in good quality images but suffers from localization bias. The second order method has better surface localization, especially in highly curved areas. It guarantees minimal detail smoothing resulting in a better visualization of polyps.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation - Viewing Algorithms; I.3.7 [Computer Graphics] Three-Dimensional Graphics and Realism

Additional Keywords: virtual colonoscopy, bias field, medical imaging, surface extraction, direct volume rendering.

1. Introduction

In recent years volume visualization has become increasingly accepted and used for many applications. Examples are microscopy, industrial inspection, simulation, games, and animations. Progress in research has resulted in increased rendering speed and high image quality.

contact address: Lorentzweg 1, 2628CJ Delft, the Netherlands, marc@ph.tn.tudelft.nl

[2]Pattern Recognition Group, Delft University of Technology

[3]Medical IT - Advanced Development - Philips Medical Systems B.V.

[4]Department of Radiology, Academic Medical Center Amsterdam

IEEE Visualization 2003,
October 19-24, 2003, Seattle, Washington, USA
0-7803-8120-3/03/$17.00 ©2003 IEEE

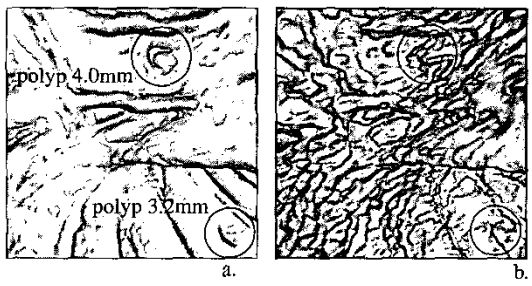

Figure 1: Virtual colonoscopic images. a) Rendering of high dose CT data. b) Rendering of simulated low dose CT data.

Obviously, the quality of visualization strongly depends on the quality of the underlying data. Noise and bias (global signal fluctuations) can significantly disturb the results. Unfortunately, high quality input data is not always available.

Virtual colonoscopy is a method to inspect the colon using a CT volume of the abdominal region (instead of physically inserting an colonoscope). As the data is acquired by CT scanning, which conventionally involves a significant radiation dose. The health risk incurred by ionizing radiation prompted researchers to try low dose CT data acquisition. In CT the image quality depends on the product of X-ray tube current (mA) and exposure time (s), usually expressed in mAs. A conventional dose is about 40 mAs, a low dose is in the order of 1.5 mAs. Alternatively, MRI could be used, which is free of ionizing radiation. This is specifically important for large-scale screening of patients.

Unfortunately, low dose CT acquisition yields noisy data. The increased noise prohibits polyp detection in images generated by direct volume rendering methods [Levoy 1988], and also by isosurface volume rendering (see Fig. 1). In addition to increased noise, MRI suffers from signal fluctuations, generally referred to as bias. Iso-surface volume rendering fails to visualize true object surfaces in such data (see Fig. 2).

Figure 2: a) MRI image, with a typical, global signal fluctuation, from the top to the bottom of the image. b) Iso-surface volume rendering at a low threshold value and c) using a slightly higher isovalue.

Much research focuses on the removal of noise effects. A solution to reduce the noise is by low-pass filtering the input data (e.g. via a Gaussian filter). However, this also results in loss of significant detail. A modification to non-linear Gaussian filtering yields proper visualization of fine details, but still requires a fairly high dose at 20 mAs [Rust et al. 2002].

Another method is to render the intensity values semi-transparently using ray casting [Kindlmann and Durkin 1998; Kniss et al. 2002], which amounts to filtering along the ray. For semi-transparent volume rendering, complex multi-dimensional transfer functions can be generated, which are based on higher-order information to differentiate between tissues and emphasize material boundaries.

In our approach, we preprocess the data using first and second order boundary detection techniques, to eliminate noise and bias. As we are mainly interested in the interior surface of the colon, a simple iso-surface rendering transfer function can be used to visualize the surface. As we do not use semi-transparent rendering, no complex transfer functions are needed.

The removal of bias fields has received much attention in recent years [Ahmed et al. 2003; Likar et al. 2000]. However, such methods can not handle noise. Additionally, there is a trade off between speed and accuracy of such a removal. A proper operation requires a considerable amount of time. We are unaware of any methods that incorporate the removal of bias in the rendering stage.

A method that solves current limitations should fulfill requirements regarding robustness against noise and a bias. Unfortunately, no existing solution meets all requirements.

This paper describes volume rendering techniques that deal with imperfections like noise and bias. The methods are based on existing first and second-order derivative filters. The novelty of our work is in the adaptation and integration of these methods for volume visualization. The algorithms are tested against the requirements above, using virtual colonoscopy data. In Section 2 the developed methods are described. Some results are presented in Section 3. We will finish by summarizing our conclusions and recommend further research (Section 4).

2. Methods

Many volume rendering techniques act by directly classifying measurement values (i.e. values within a certain range are rendered translucently, the others opaquely). We propose two methods that are based on first and second order derivative information respectively. Henceforth, we will refer to these methods as first-order and second-order solutions described in sections 2.2 and 2.3. The first order solution is inspired by the Canny edge detector, which is introduced in section 2.1.

2.1 Canny Edge Detection

A common first order edge detector was developed by Canny [1987] (Canny Edge Detection, CED). Initially it was applied to 2D images. CED is based on the gradient magnitude. The algorithm comprises three steps:

- *Gradient calculation*
  The gradient vector is calculated at each position. A method for this is to convolve the data with derivative of the Gaussian filters in all three directions (x,y,z).

- *Non-maxima suppression*
  Positions that are not local maxima in the direction of the gradient vector are excluded. The operation yields a surface only one voxel thick at which the magnitude is preserved [Canny, 1987].

a. b.

Figure 3: Canny Edge Detection. a) Input image (T1 weighted MRI data). b) Output image.

- *Hysteresis thresholding*
  A double threshold is used on the gradient magnitude to detect the true surface and to close holes. Magnitude values above the 'high' threshold are immediately assigned the status of surface voxel. This threshold is set higher than the noise level. Thus, spurious surface fragments are removed. A voxel above the 'low' threshold is assigned the status surface voxel if one of its neighbors already is one. Iterative application of this rule yields the surface.

CED is fast and robust against noise. However, global intensity fluctuations will result in incorrect surfaces. In high intensity areas spurious edges are detected and in low intensity areas the surface becomes discontinuous (see Fig. 3).

We have adapted the original Canny edge detector by applying it in a truly 3D manner, and using a locally normalized edge indicator, as explained below.

2.2 The first order solution

The volume is visualized by a modified ray casting technique. The preprocessing consists of three stages (compare with Section 2.1):

- *Gradient calculation*
  The gradient vector is calculated at each position (A) using Gaussian derivative filters.
  The local gradient magnitude ($\|G_A\|$) is divided by the local maximum of the gradient magnitude ($\|G_{max}\|$), which is determined on the line through A that is parallel to the gradient vector. The search window around A extends 2.5 mm in both directions which corresponds to the size of the smallest relevant detail. By definition the values of the normalized gradient magnitudes vary between zero and one. The responses are raised to a power p ($p > 1$) for sharpening:

$$E = \left(\frac{\|G_A\|}{\|G_{max}\|}\right)^p \qquad (1)$$

  An example image is shown in Fig. 4.

- *Non-maxima suppression*
  By definition only the responses exactly equal to one represent local maxima in the gradient direction. This is because of the local normalization. However, such a strict criterion yields a severely fragmented surface (due to sampling artifacts, see below). To compensate for such artefacts, the positions eith sample values that are smaller than a threshold $t<1$ are excluded. Consequently, the method becomes less sensitive to such errors.

• *Thresholding*
A threshold is used on the non-normalized gradient magnitude to include only true surface positions. This threshold is set higher than the noise level to remove spurious surface fragments.

The algorithm renders the included postions on a ray before the first local maximum with an opacity set to zero and after it with an opacity set to one.

We must be careful with the calculation of surface normals. Consider using Gaussian first derivative filters. The surface to be visualized is a thin sheet in 3D (one can visualize it as a ridge in 2D, see Fig. 4b). The gradient on the top of the ridge approximates the zero-vector. Consequently, the derivatives become unstable near to the top.

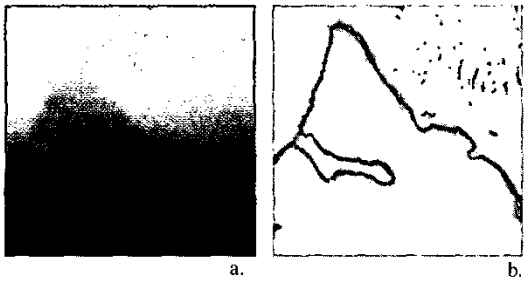

a. b.

Figure 4: a) MRI volume (notice the bias from top to bottom) b) Normalized Gradient volume.

A more robust estimation of the normals is obtained using the Structure Tensor (also referred to as the Gradient Square Tensor) [Kass and Witkin 1987]. The Structure Tensor (F) is defined by the dyadic product of the gradient vector with its transpose:

$$\bar{F} = \overline{(\nabla f)(\nabla f)^t} = \begin{bmatrix} \overline{f_x f_x} & \overline{f_x f_y} & \overline{f_x f_z} \\ \overline{f_x f_y} & \overline{f_y f_y} & \overline{f_y f_z} \\ \overline{f_x f_z} & \overline{f_y f_z} & \overline{f_z f_z} \end{bmatrix} \quad (2)$$

f_x, f_y and f_z denote the partial derivatives of the data $f(x,y,z)$, which we calculate by a convolution with a Gaussian derivative filter. Another Gaussian is used to regularize the Structure Tensor (depicted by the overhead bar). Fig. 5 shows a visual interpretation of the Structure Tensor.

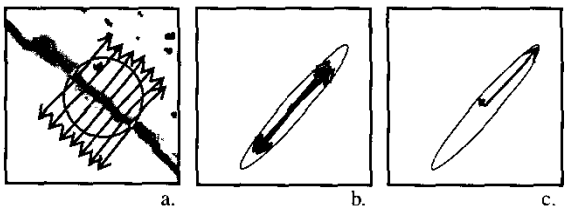

a. b. c.

Figure 5: The Structure Tensor in 2D. a) A ridge with opposite gradient directions. b) Gradient magnitude vectors are assigned to one point. The principal axes are found using an eigensystem analysis of the tensor data. c) The largest eigenvector of the tensor is perpendicular to the ridge.

It must be emphasized that the Structure Tensor is only used to determine the orientation of the surface normals. The location of the surface itself is not affected. Also, the normals must be calculated from the normalized gradient volume and not from the input data. Due to the normalization step, the normals in both spaces may not have the same direction.

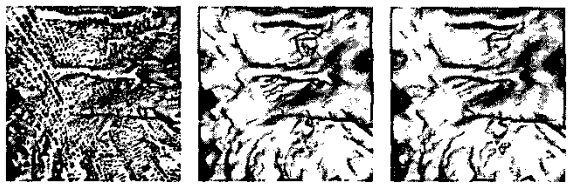

a. b. c.

Figure 6: Example renderings at t=0.9 (a), t=0.7 (b), t=0.5 (c).

The first order solution has four (sets of) critical parameters:

- the sharpening factor p
- the threshold t
- the width of the Gaussian first derivative filters $\sigma_{x,y,z}$
- the degree of tensor smoothing σ_T

A larger p sharpens the response and makes details on the surface appear more jagged. A small p smooths noise, but at the expense of detail. Example renderings at t=0.9, t=0.7, t=0.5 are given in Fig. 6 ($\sigma_{x,y,z}$ = 1.0mm, σ_T = 1.5mm). Clearly, if the isovalue t approximates one (which corresponds to the true edge), the result suffers from severe sampling artifacts.

$\sigma_{x,y,z}$ is used to suppress the effect of spatially uncorrelated noise on the derivatives. Conventionally, it is chosen small, to correspond as much as possible to the scale of the smallest relevant detail (i.e. the smallest features of interest). σ_T is tuned to minimize the influence of normals from spurious detail.

By using derivatives the method becomes sensitive to noise. This effect is compensated for by the width of the derivatives. A clear advantage of the method is that it does not depend on an isovalue of absolute data values.

The effect of the parameter settings on the rendering will be explored in Section 3.2.

2.3 The Second Order Solution

Another method to identify the edge is via the zero crossing of a second derivative operator. Common second order edge detectors are the Marr-Hildreth detector, defined as the Laplacian of Gaussians (LoG), and the second derivative component in the gradient direction (SDGD) [Marr and Hildreth 1980]:

$$LoG = f_{xx} + f_{yy} + f_{zz}$$
$$SDGD = f_{gg} \quad (3)$$

in which $f_{xx}, f_{yy} f_{zz}$ and f_{gg} are second derivatives of the image function f calculated via Gaussian derivatives; x,y,z are the main coordinate directions and g is the direction of the gradient vector. Although these detectors are widely applied in 2D, we use a truly volumetric 3D generalization.

With curved object boundaries there is always a trade-off between noise suppression and surface displacement . A larger filter size reduces noise, but at the expense of increased surface displacement. First order derivative filters, such as the Gaussian

derivative filter of the previous section, cause a surface displacement that is always in the direction of positive curvature (as shown in the fig. 7). Unfortunately, it is generally impossible to compensate for the first derivative surface displacement.

The LoG and SDGD filters have opposite surface displacements. The operators may be combined [Verbeek and Van Vliet 1994] to reduce the effect, as the opposite-signed displacements cancel each other to a certain extent. The combined edge detector is generally referred to as the *Plus operator*:

$$Plus(f) = SDGD(f) + LoG(f)$$

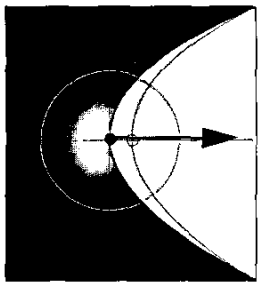

Figure 7: Filtering a curved object yields a dislocated surface. The filter neighborhood (depicted by the circle) covers a larger area outside the object than inside. Consequently, a dislocated edge is detected in the direction of positive curvature (indicated by the arrow).

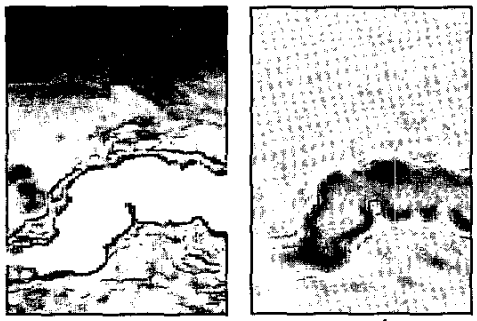

a. b.

Figure 8: The left image shows the input data, the right image the result of the Plus operator.

Fig. 8a presents MRI data with contrast fluid injected into the colon appearing as highly intense. The application of the Plus operator results in a new volume (Fig. 8b) which is referred to as the *Plus volume*. A transition visible is in Fig. 8b from negative values (black) to positive values (white). In between is the zero crossing.

The Plus volume is visualized by a modified ray casting algorithm. Initially, those voxels are selected that have a gradient magnitude higher than a certain threshold. Thus, a coarse region of interest is identified. This threshold is applied to reduce noise. We have found that the value of the threshold is not very critical for the result. Subsequently, the algorithm renders the samples on a ray with an opacity set to zero before the first zero crossing and with an opacity set to one after it.

The surface normals are calculated by Gaussian first derivative filters in the Plus volume. This is possible because the gradient vectors around the transition do not have opposite directions.

The second order solution has two sets of critical parameters:

- the width of the Gaussian second derivative filters. (Equation 4)
- the width of the Gaussian first derivative filters σ_n (used to calculate surface normals).

Only the widths of the second derivatives affect the localization of the surface. The second derivative parameter balances noise against surface detail. σ_n merely affects the shading. So, σ_n should be as small as possible for obtaining stable normals. The effect of the parameter settings will be explored in Section 3.

3. Results

The first and second order solutions were tested both visually and quantitatively. The implementation we use preprocesses both the normalized gradient volume (Fig. 4) and the Plus volume (Fig. 8). This results in an fully automatic preprocessing step and a near real-time ray casting.

3.1 Data acquisition

A CT volume was randomly selected from an ongoing virtual colonoscopy study. The original CT scans were acquired at a level of 40 mAs. To have a reference, the low dose CT scans were generated from these normal dose scans.

Low-dose CT scans were reconstructed at 1.5 mAs by modifying the raw transmission data of each spiral CT scan using a simulation technique. The raw transmission measurements were modified by adding a random number from a normal uncorrelated distribution with zero mean and a variance according to the desired simulated dose setting. A simulated 50 mAs scan, for example, was obtained from the 100 mAs scan by adding a random number with the same variance as the raw transmission measurement. Because noise is added before 3D data reconstruction, the simulation is valid [Mayo et al. 1997]. Typically, the volume consisted of 512x512x256 voxels at a resolution of 0.6x0.6x1.2 mm^3. Example images are shown in Figure 9 (the noise is correlated).

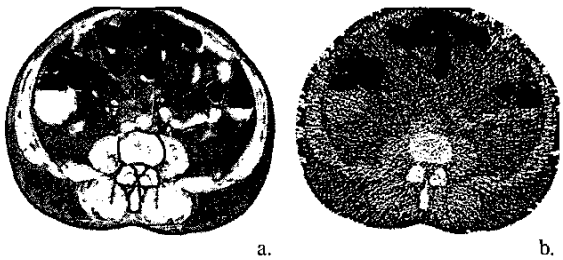

a. b.

Figure 9: Slice of CT data. a) Directly from scan. b) Simulation of low dose.

An MRI volume was randomly selected from an ongoing study. The data was acquired via a Siemens 1.5 Tesla scanner. A typical volume consisted of 256x256x128 voxels sized 1.2x1.2x2.0 mm^3.

The first and second order solutions were implemented on an experimentally enhanced version of the Philips EasyVision workstation.

3.2 Visual Results

The 3D visualizations of low dose CT data (Fig. 1, Fig. 6, Fig. 11-Fig. 10 and Fig. 16) all show the same scene with the same viewing parameters.

Figure 10: a) Iso-surface volume rendering of high dose CT (repeated from Fig. 1a). b) First order method on low dose CT. c) Second order method on low dose CT.

An optimum first order rendering ($\sigma_{x,y,z}$=1.2 mm, σ_t = 1.0 mm, p = 8) is shown in Fig. 10 b and the optimum second order rendering is shown in Fig. 10 c ($\sigma_{x,y,z}$ = 1.2 mm, σ_n = 1.0 mm). The parameter settings will be discussed below. Notice that the first and second order images are rendered with the same filter sizes. Fig. 10 a. shows an iso-surface volume rendering of the high dose CT volume. We consider this the reference standard.

The effects of significant parameter adjustments on the first order rendering are depicted in Fig. 11. Each row shows the result as a single parameter is varied (from top to bottom: $\sigma_{x,y,z}$, σ_t, p respectively). The effect of the parameters in the second order solution are depicted in Fig. 12. The optimum parameter setting (cf Fig. 10b and c) was determined based on visual inspection of these renderings.

To study the effect of bias fields we selected two camera positions in an MRI volume, one with and one without visible noise and signal fluctuations. Fig. 13 shows the result for large signal fluctuations and noise, Fig. 14 for the other with a small bias. The parameter settings were again found experimentally. Both images were rendered with the same parameter settings. As expected, the iso-surface volume rendering shows holes and surface fragments, due to the signal fluctuations (Fig. 14).

Figure 13: Result on MRI in area of high contrast and low fluctuation. a) Iso-surface volume rendering. b) Ray casting of normalized gradient volume results in smoothing. c) Plus operator.

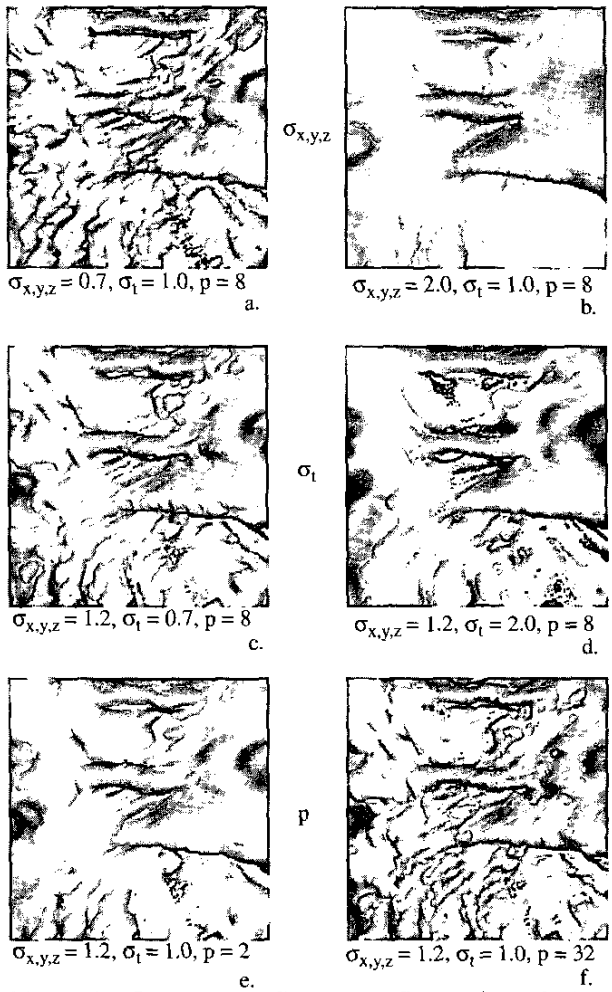

Figure 11: Parameter influence on first order solution. (all sigma's in mm) a) Low $\sigma_{x,y,z}$, resulting in a noisy surface. b) High $\sigma_{x,y,z}$, resulting in a smooth surface at the expense of detail. c) Lower σ_t. d) High σ_t. e) Low p. f) High p.

3.3 Localization Accuracy

A reference image for the CT data was defined by iso-surface volume rendering of the (normal) 40 mAs data. The threshold for the iso-surface was determined by maximizing the sum of gradient magnitudes over all surface positions. The surface displacement was defined by the distance along viewing rays between the reference image and the first and second order surface respectively. We have opted for this approach to illustrate the localization accuracy from a typical view point. A reference image cannot be

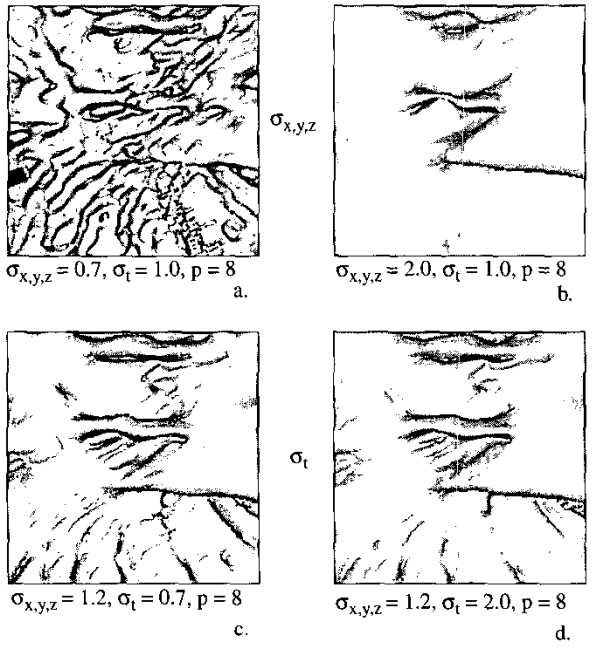

Figure 12: Parameter influence on second order solution. (all sigma's in mm) a) Low $\sigma_{x,y,z}$, resulting in a noisy surface. b) High $\sigma_{x,y,z}$, resulting in a very smooth surface at the expense of detail. c) Lower σ_n. d) High σ_n.

Figure 14: Result on MRI in area of low contrast and high fluctuation. a) Iso-surface volume rendering. b) Normalized Gradient Magnitude. c) Plus operator.

easily defined for MRI data due to lack of a proper standard. For that reason we have restricted the quantitative analysis to CT data.

A small percentage (approximately 0.5%) of very large errors was caused by small variations in fold thickness (see Fig. 15). Such artifacts were eliminated from the analysis by removing the errors larger than 10 mm.

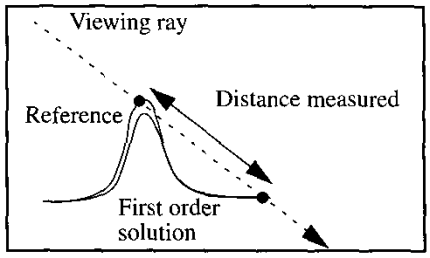

Figure 15: Example of a measurement artifact in the surface displacement.

The mean first order error was 0.6 mm towards the camera (averaged over ten randomly selected positions). The errors were within 3 mm in 95% of the cases. The second order solution had a mean localization error of 0.19 mm (again averaged over the same ten positions); the errors were below 1.0 mm in 95% of the cases .

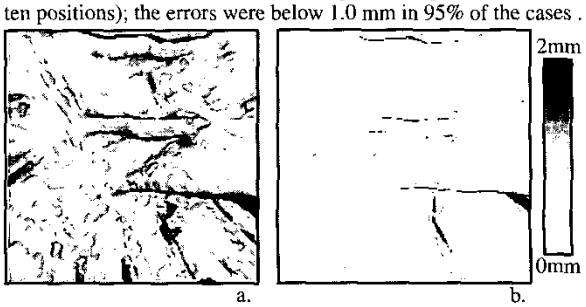

Figure 16: The distribution of the surface displacement shown as a greyscale image. a) First order solution. b) Second order solution.

Fig. 16 depicts the surface displacement for both solutions. Obviously, the first order solution has a significantly larger localization error (Fig. 16a) than the second order solution (Fig. 16b).

3.4 Performance

The rendering time is divided into preprocessing time and time needed for ray casting. The preprocessing for the first order solution took two hours (which includes calculating the Structure Tensor for the whole volume) and the ray casting only took on the order of a few seconds per image. The whole data set was preprocessed for the second order solution in 40 minutes and the ray casting took about one second. These numbers were obtained for the CT data using an 2.0 GHz computer with 256 MB of internal memory. For other data volumes the preprocessing time is scaled in proportion with volume size.

4. Conclusions

This paper reports on progress in volume visualization of data with noise and bias fields. We presented methods for volume rendering based on first and second order derivative information.

First, a modified version of the Canny edge detector was applied, to suit application in virtual colonoscopy. The modification included a method for robust surface normal estimation. This first order technique effectively deals with noise, but introduces a surface displacement.

Next, a second order visualization method was presented incorporating the Plus operator that combines the Marr-Hildreth operator and the Laplacian of Gaussians. This operator yields accurate surface localization, and it is insensitive to global signal fluctuations. The visual results show good surface quality retaining important shape details.

The methods were tested on simulated low-dose CT data (derived from regular-dose data). The new techniques were compared to iso-surface volume renderings from the regular dose data (low noise, no global fluctuations). The comparisons showed good agreement in visualization of significant details, such as polyps diagnosed by radiologists.

Currently, we work on a more rigorous test. We will develop a more accurate and view independent measure for the surface localization. We also will perform extensive clinical tests. At last, we intend to optimize the methods for speed.

The results presented suggest that low-dose CT and MRI may become feasible as techniques for virtual colonoscopy in the future. This may be a step towards the use of virtual colonoscopy in large-scale clinical screening for early diagnosis of colonic polyps.

5. Acknowledgments

The authors would like to thank dr. F.A. Gerritsen of Philips Medical Systems for his support. The CT and MRI data was provided by the Academical Medical Center Amsterdam. We would like to thank: Henk Venema, Rogier van Gelder and Jasper Florie for their input.

References

AHMED, M.N., YAMANY, S.M., MOHAMED, N., FARAG, A.A., AND MORIARTY, T., 2003. A Modified C-Means Algorithm for Bias Field Estimation and Segmentation of MRI Data. *IEEE Transactions on Medical Imaging 21*, 193-200.

CANNY, J.F., 1986. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence 8*, 679-698, 1986.

KASS, M.. AND WITKIN, A., 1987. Analyzing Oriented Patterns. *Computer Vision Graphics and Image Processing 37*, 362-385.

KINDLMANN, G., AND DURKIN, J.W., 1998. Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering. In *Proceedings of the IEEE Symposium on Volume Visualization*, 79-86.

KNISS, J., KINDLMANN, G., AND HANSEN, C., 2002. Multi-Dimensional Transfer Functions for Interactive Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics 8*, 270-285.

LEVOY, M., 1988. Display of Surfaces From Volume Data. *IEEE Computer Graphics and Applications 8*, 3, 29-37.

LIKAR, B., VIERGEVER, M.A., AND PERNUS, F., 2000. Retrospective Correction of MR Intensity Inhomogeneity by Information Minimization. In *Proceedings of MICCAI 2000*, 375-384.

MARR, D., AND HILDRETH, E.C., 1980. Theory of Edge Detection. In *Proceedings of the Royal Society 207*, 187-217.

MAYO, J.R., WHITTALL, K.P., LEUNG, A.N., HARTMAN, T.E., PARK, C.S., PRIMACK, S.L., CHAMBERS, G.K., LIMKEMAN, M.K., TOTH, T.L., AND FOX, S.H., 1997. Simulated Dose Reduction in Conventional Chest CT. *Radiology 202*, 453-457.

RUST, G.F., AURICH, V., AND REISER, M., 2002. Noise/Dose Reduction and Image Improvements in Screening Virtual Colonoscopy with Tube Currents of 20 mAs with Nonlinear Gaussian Chains. In *Proceedings of SPIE*, vol. 4683, 2002.

VERBEEK, P.W., AND VAN VLIET, L.J., 1994. On the Localization Error of Curved Edges in Low-Pass Filtered 2-D and 3-D Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*, 726-733.