

Hierarchical Volume Analysis and Visualization Based on Morphological Operators

Christoph Lürig and Thomas Ertl

Computer Graphics Group
Universität Erlangen-Nürnberg, Germany*

ABSTRACT

One common problem in the practical application of volume visualization is the proper choice of transfer functions in order to color different parts of the volume meaningfully. This interactive process can be very complicated and time consuming. An alternative to the adjustment of transfer functions is the application of segmentation algorithms. These algorithms are often dedicated to a limited range of data sets and tend to be very compute intensive.

In this paper we propose a morphology based hierarchical analysis to estimate the optical properties of the volume to be rendered. This approach requires fewer parameters and incorporates also spatial information, but it is far less compute intensive than most of the segmentation methods. The hierarchical analysis is constructed in analogy to the wavelet analysis, except for the fact, that non-linear filters are used in our case. These morphological operators have a lower distortional influence on the analyzed structures than the usual linear filters.

A special decomposition of the morphological operators will be discussed, that leads to an efficient implementation of this approach. This technique reduces the three dimensional analysis to a one dimensional computation, as it is done in tensor product based linear filters. The resulting decomposition may also be parallelized easily. We demonstrate the usefulness of the proposed technique by applying it to medical and technical data sets.

1 Introduction

Direct volume rendering has become a widely applied visualization method over the last several years. An overview on volume visualization in general can be found in the book of Kaufmann [4]. One of the major practical problems in volume visualization is the proper choice of transfer-functions in order to map the data to meaningful optical properties of the volume to be visualized. This is often problematic as the application of transfer-functions just maps from the defined function value to these optical properties and does not incorporate spatial information.

One alternative to the application of transfer-functions are segmentation methods, but those are usually application specific and quite time consuming. Segmentation methods are built to generate an initial symbolic description of a data set. The description may for instance be a labelization or the representation of an edge, but either way, it is always associated with a discrete decision. Many segmentation methods arise from pattern recognition problems and have been constructed to identify distinct objects. A general overview of basic segmentation methods can be found in the book by Niemann [11].

In contrast to traditional segmentation methods, the method we propose does not generate discrete decisions. This means that the

uncertainty associated with the proposed transformation is also reflected in the final visualization of the data set. The method we propose is a multi-scale method, which is based on morphological operations. This approach requires fewer parameters to be adjusted by the user and incorporates spatial information. In the final result, regions of the volume, where a defined frequency is present, are labeled with a certain color. This multi-scale method has been constructed in analogy to the wavelet transformation.

The application of scale-space theory has become very popular in the computer graphics and pattern analysis community along with the increasing application of the wavelet transformation. The wavelet transformation has not only been used for image compression as described by Stollnitz et. al. [16], but also to detect edges and to characterize signal properties as explained in the articles of Mallat et. al. [8, 6, 7]. Gauch and Pizer detect certain image features by tracking signal features through scale space [1]. Westermann applied the idea of the wavelet transformation in [18] to the area of volume visualization. Here again, not only the possibility of data compression is important, but also the idea of edge information that is represented by the wavelet coefficients. Westermann renders the wavelet coefficients, that represent the gradient information at a certain level of detail [19]. The idea to incorporate gradient information into the volume rendering process was first presented by Levoy [5]. In contrast to the wavelet approaches he does not make use of multi-scale information. Muraki [9] applied the wavelet transformation to volume data sets, to render the Laplacian of the Gaussian of these data sets.

The disadvantage of the wavelet transform and the other scale-space methods, that are based on linear filters, is the fact that the analyzed structures become distorted, especially on lower frequency scales. One class of special non-linear filters, that do not flatten the original contours are the morphological operations. Morphological operations compute the minimum and the maximum of pixel values within a certain scope. Before the maximum and the minimum is computed, values of the so called structuring element may be added or subtracted. A good introduction to gray-scale morphology can be found in Sternberg [15]. The operation, that is applied in this work, is the opening operator, which is a combination of the erosion and dilatation operator.

The morphological operators have been mainly used in pattern recognition to perform some kind of preprocessing before segmentation methods are applied. Very often, the morphological operations are used in combination with region growing methods. This combination is applied by Höhne and Hanson [2] and by Saiviroonporn et. al. [12]. They both use the idea to separate distinct organs in tomographic data sets. First, a volume of interest is defined by the user, then an opening operator is applied to break down the small bridges connecting the organs, and finally, the organs themselves are identified by applying a region growing algorithm to the prepared data set. This is in contrast to the approach presented here, where we do not want to generate surface descriptions of distinct organs or other volumetric objects. Our aim is to color different regions according to the spatial frequency, that is present within them and to make use of the vagueness, that can be expressed by volume

*Lehrstuhl für Graphische Datenverarbeitung (IMMD9), Universität Erlangen-Nürnberg, Am Weichselgarten 9, 91058 Erlangen, Germany, Email: cpluerig, ertl@informatik.uni-erlangen.de

rendering. This vagueness is caused by unsharp borders, that are generated by this soft segmentation method.

Most of the work, that combines morphological operators with scale-space theory has been done in order to solve a certain class of partial differential equations. Such examples are described in van den Boomgard and Smeulders [17] and in Sapiro et. al. [14]. Both papers rely on erosion or dilatation operators of increasing size to generate their scale space. The diamond structuring element, that is used in our algorithm, is also introduced there. Jackway and Deriche [3] use the scale-space approach to track image features across different scales, in order to see when they disappear. In contrast to these approaches, that use either erosion or dilatation operators, we use the opening operator for each scale. The resulting scale is computed by the difference of the original image and the opened one. This is done in analogy to the wavelet transformation. The opening operator suppresses certain details of the original data set. The result of the opening operator corresponds to the smoothing coefficients in wavelet decomposition. Thus the difference between the original and the opened data set corresponds to the wavelet coefficients and represents a certain frequency spectrum. The resulting coefficients are then mapped to the color emission values of the volume directly using a threshold and a constant multiplication factor.

Sakas and Walter [13] have constructed a multi-resolutional analysis, that makes use of morphological operators by applying an opening operation to the low pass filtered data set. In contrast to our method that uses opening operators of different size, the multi-scale aspect is provided by the usage of different frequencies for the low pass filter. The opening operator is kept the same. Also in contrast to our approach they apply binary morphology to the thresholded data set.

The major problem in the application of morphology is, that morphology is still a very time consuming operation. In order to tackle this problem, we decompose the expensive operation with a large structuring element into several smaller ones as it is explained in Zhuang and Haralick [20]. They reduce the size of the structuring element, but keep its dimension. In this paper we will demonstrate a technique to decompose the multi-dimensional structuring elements into effectively one-dimensional ones using a pseudo tensor product approach. The resulting morphological operations can be efficiently parallelized using standard POSIX threads (see Nichols et. al. [10]).

In the next section we describe the data analysis process in more detail. The separation and the processing of the different frequency spectra is discussed here. Special emphasis is put on the morphological operators, the scale-space aspect and the analogy to the wavelet analysis. Section 3 focusses on the efficient implementation of the morphological operators. It discusses the decomposition of the operators and the pseudo tensor product based approach including a comparison of the complexity, as well as the possibilities of parallelization. Some experimental results including medical and technical applications together with performance evaluations are presented in section 4. We conclude with some ideas on extending the method and describing possibilities for future work.

2 Hierarchical Analysis

Morphological operators are a special form of non-linear filters, that have been mainly invented to separate or to combine different objects in an image with a minimum distortion of the contours. Primarily, morphological operations are defined for binary images and are then extended to gray-scale images. Consider X and Y to be two binary images. Then the erosion $Z := X \ominus Y$ is defined as

$$Z_{(i,j)} = \begin{cases} 1 : \forall(x,y) : Y_{(x,y)} = 1 \rightarrow X_{(i-x,j-y)} = 1 \\ 0 : otherwise \end{cases} \quad (1)$$

The second of the basic morphological operators is the dilation operator $Z := X \oplus Y$ which can be considered as the dual operator to the erosion:

$$Z_{(i,j)} = \begin{cases} 1 : \exists(x,y) : Y_{(x,y)} = 1 \wedge X_{(i+x,j+y)} = 1 \\ 0 : otherwise \end{cases} \quad (2)$$

The erosion operator cuts away certain parts of the boundary of the analyzed image. The amount that is cut away is specified by the structuring element Y , which is usually quite smaller than the image. In contrast the dilation operator enlarges the analyzed image structures.

These basic operators can then be combined to more complex operations. The next two operations that evolve from erosion and dilation are the opening and the closing operator. The closing operator is defined as

$$X \circ Y = (X \oplus Y) \ominus Y \quad (3)$$

and the opening operator is defined as

$$X \circlearrowleft Y = (X \ominus Y) \oplus Y \quad (4)$$

Generally speaking the opening operator tends to break up small bridges between connected image regions whereas the dilation operator fills small gaps in solid components. In this paper we apply the opening operator in order to analyze the data sets.

These operators are just defined on binary data up to now. In order to extend the definition to gray-scale data, the gray-scale data is transformed to a binary data set with an extra dimension. In this way an image is transformed to a three-dimensional data set. The axis pointing along the dimension that corresponds to the gray-scale is filled with elements up to the height corresponding to the gray level value. In order to implement gray-scale morphology, this lifting into the extra dimension can be done implicitly by appropriately defining the dilation and the erosion operators for a two-dimensional data set in gray-scale. The dilation operation $Z := X \oplus Y$ with the structuring element Y is now defined as

$$Z_{(i,j)} = \max_{x,y} \{X_{(i-x,j-y)} + Y_{(x,y)}\} \quad (5)$$

where Z , X and Y represent the gray-scale values of the data set. In accordance the erosion operator $Z := X \ominus Y$ becomes

$$Z_{(i,j)} = \min_{x,y} \{X_{(i+x),(j+y)} - Y_{(x,y)}\} \quad (6)$$

The opening and closing operators are defined automatically for gray-scale morphology as they are constructed again from a combination of the erosion and dilation operators. We use the opening operator for two dimensions on images and for three dimensions for volume rendering purposes. The three-dimensional definition is analogous to the two dimensional case. All these definitions can be found in Sternberg and Stanley [15].

The structuring element we use exhibits a diamond structure. Depending on the size of the element, which is indicated by max where the indices range from $0 \dots max$, and an additional scaling factor, which is defined by $factor$, the structuring element Y is defined as

$$Y_{i,j} = factor \cdot \left(\cdot max - \left| i - \frac{max}{2} \right| - \left| j - \frac{max}{2} \right| \right) \quad (7)$$

for the two dimensional case. As an example, the structuring element with the values $factor = 1$ and $max = 2$ is sketched in Fig. 5.

The basic transformation, that is applied in the presented algorithm is the opening of the data set and the computation of the difference to the data set, that has been opened with a smaller structuring element. For the case, that this smaller element has an extent

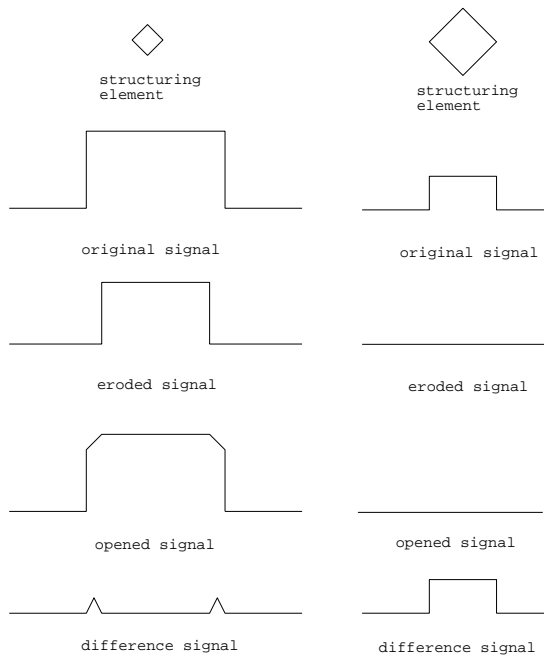


Figure 1: The result of the analysis procedure

of one ($max = 0$), which is a neutral element with respect to the opening operation, this process is illustrated for a one-dimensional gray-scale signal in Fig. 1. On the left hand side a coarse structure is shown, where in fact the boundaries are detected. The right hand side shows a smaller structure, that is eliminated in total by the opening operation with a larger structuring element, and thus the difference signal reflects the complete original signal.

In order to generate a multi-scale method, the differences are computed with structuring elements of increasing size. The resulting steps correspond to different spatial frequency spectra. This construction has been done in analogy to the wavelet transformation, where an original signal is low passed filtered several times using some smoothing function. The computed wavelet coefficients for different scales correspond to the computed difference data sets in our case. In contrast to the wavelet transformation, the increase in size of the structuring element can vary from scale to scale. The morphological operations are not injective and thus not invertible, which is a remarkable difference to the wavelet transformation. In the case of morphological operations there are also no orthogonal sub-spaces, which are constructed.

In principle, the morphological analysis could be applied for many different scales. For practical reasons, however, we restrict the algorithm to four scales resulting in three difference signals which correspond to low, middle and high spatial frequencies. For a three-dimensional scalar data set this means, that we compute three quantities which are the intensities of the respective frequency band for each voxel. In order to convey this information we use direct volume rendering which allows us to visualize structures in a volumetric data set by assigning different colors and opacities. In contrast to traditional volume ray-casting, where the RGBA values are derived directly from the scalar values of the data set by means of transfer-functions, we generate the three color components red, green and blue and the absorption coefficient directly from the computed intensities of the frequency spectra, thus incorporating also spatial information.

The overall data flow of the analysis process is shown in Fig. 2. Four opening operations with structuring elements of increasing

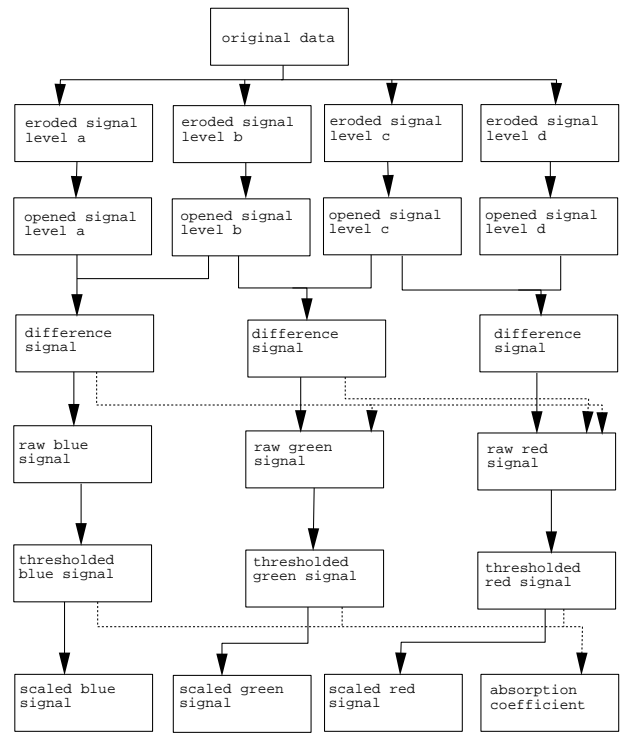


Figure 2: The overall data flow of the morphological analyzer

size are performed all together. The sizes of the applied structuring elements should be about the feature size, that is to be colored. The three color components result from the differences between two different scales each. The blue component evolves from the difference of the second and the first opening operation, the green from the second and the third one and the red component from the fourth and the third opening operation. In order to generate the final color values, the resulting three components are assigned priorities among each other and they are thresholded with an indicated absolute value. The global thresholding has been proven to be useful for suppressing artifacts and noise in the data set. If the signal of a high frequency band exceeds a certain threshold, the signals for the lower frequency bands are suppressed. The priority of the higher frequency bands is reasonable, since several frequency bands quite often occur in the same position. The signals of all bands are normalized to a given value range and may be multiplied by a constant factor later on. This frequency information can also be interpreted as a characterization of the feature size. The opacity is derived from a linear combination of the three emission values for the three colors. The resulting RGBA data are finally fed to a simple volume ray-caster, that uses these values directly for the computation of the light transport equation. A general impression of this method can be seen in image 7(g). Here, the multi-scale morphological analysis of the CT-scan of a cup automatically results in a color coding of the wall thickness.

Since some of the features of the proposed method are difficult to understand in the three-dimensional case, we also provide some two-dimensional synthetic examples. In these examples, the three frequency scales which are actually processed are mapped to a gray-scale value. An absorption coefficient is not computed. Since all frequency spectra are indicated by the same range of gray-scale values, we have provided an additional annotation.

A first example is shown in Fig. 3 where the original image in the upper left corner consists of three squares of different size. The

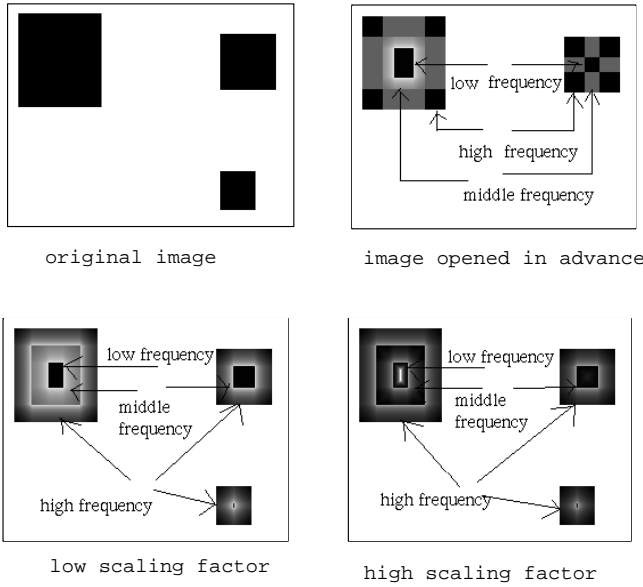


Figure 3: Two-dimensional example for the procedure

image shown in the lower left corner represents an analyzed data set, where no initial opening has been performed. This means that the structuring element used at level a has an extension of one. The big square consists of three regions of different sizes. The central region is due to the lowest frequency part. Towards the outside of the rectangles, the regions that correspond to the middle and the high frequency band can be found. The regions are separated quite strictly, as the lower frequency bands are suppressed, if there is a signal with a minimum strength at one of the higher frequency bands.

The image in the lower right corner has been computed with a structuring element of the same size, but the value *factor* of equation (7) has been increased. This results in a hole in the middle of the big rectangle. The extra scaling factor increases the pitch of the diamond and thus reduces the area where a non-zero difference appears.

Finally the rectangle in the upper right image has been generated by a filter cascade, where the first opening operator, which has been applied at level a , has used a relatively large structuring element. The result is, that this opening operation completely eliminated the small rectangle. The appearance of the other rectangles has been left nearly unchanged. This could have not been accomplished using linear filter techniques, since the small rectangle would never disappear entirely and the structures of the larger rectangles would become distorted. The approach of Sakas and Walter [13] would also produce this effect, as the applied low-pass filter in combination with a small morphological operation would not be able to suppress the small rectangle.

One can also see that there is a preference of horizontal and vertical edges especially in the image in the upper right corner. The anisotropy phenomenon is due to the diamond structure of the structuring element. This anisotropy phenomenon can also be seen in Fig. 4, where a number of circles have been transformed. Especially in the representation of the lower frequency band, one can clearly see the preference for vertical and horizontal edges. The reason for choosing this structuring element will become clear in the next section.

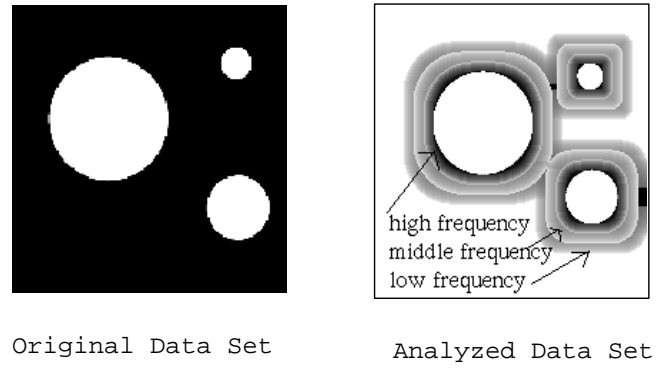


Figure 4: Demonstration of the anisotropy property of the applied structuring kernel

3 Efficient Algorithm

The general problem, that comes along with the application of morphological operations is the efficiency aspect. The computational complexity increases with the square of the size of the structuring element in the two-dimensional case and with cubic complexity in the three-dimensional case. This is especially problematic in our case, where structuring elements of sizes up to 17 are used to identify structures in the analyzed data set.

3.1 Structuring Element Decomposition

The general approach to cope with the efficiency problem is to decompose a large structuring element into several smaller ones. The basic relations, that are applied here are described in Zhuang and Haralick as [20]

$$X \ominus (Y \oplus Z) = (X \ominus Y) \ominus Z \quad (8)$$

and respectively

$$X \oplus (Y \oplus Z) = (X \oplus Y) \oplus Z . \quad (9)$$

This means, if one can find a decomposition of a large structuring element S into several smaller ones:

$$S = H_1 \oplus H_2 \oplus H_3 \oplus \dots \oplus H_N , \quad (10)$$

one can implement the morphological operation with a large structuring element by the application of several smaller ones:

$$X \ominus S = (((X \ominus H_1) \ominus H_2) \dots) \ominus H_N \quad (11)$$

$$X \oplus S = (((X \oplus H_1) \oplus H_2) \dots) \oplus H_N . \quad (12)$$

This technique is also applied in linear filtering, where for instance the convolution with a large Gaussian can be simulated by the iterated convolution with several small Gaussian filters. The diamond structuring element, that is used in our case can be simulated by iteratively applying the structuring element shown in Fig. 5 for the two-dimensional case. When this structuring element is dilated by itself, with the values outside the shown 3x3-scheme assumed to be $-\infty$, one gets a new pyramidal element with a new *max* size of four. For each further iteration, the structuring element size increases by another two elements in each dimension. The extra scaling factor *factor* can also be applied to the small structuring elements and it is propagated through the iteration process. For a structuring element to keep its shape when being self-dilated iteratively a necessary condition is, that it is convex and symmetrical

0	1	0
1	2	1
0	1	0

Figure 5: The primitive structuring element for two dimensions

with respect to all coordinate axes. These properties also hold for the diamond structuring element, which is applied here.

If we assume the extension of a three-dimensional structuring element to be $2n + 1$, then the number of additions and comparisons to be performed for the computation of the value for one voxel sum up to $8n^3 + 12n^2 + 6n + 1$. For the iterative application of the smaller structuring element the overall complexity reduces down to $27n$. The iterative application of structuring elements results in piecewise linear structuring elements in consequence. Using this kind of small sized atomic structuring element, we get the already discussed diamond structure.

In order to simulate the structuring element indicated in equation 7, we propose a decomposition into three elements for the three-dimensional case. The first of these structuring elements is defined as follows:

$$Y_a(i, j, k) = \begin{cases} factor \cdot (\frac{max}{2} - |i - \frac{max}{2}|) & : j = 0 \wedge k = 0 \\ -\infty & : otherwise \end{cases}$$

The structuring elements Y_b and Y_c are defined analogously for the y- and the z-axis, which results in a permutation of the i, j, k indices. These elements have entries of $-\infty$ at most positions except for one single string parallel to one of the coordinate axis. The final structuring element can be composed of the three structuring elements:

$$Y = (Y_a \oplus Y_b) \oplus Y_c \quad (13)$$

The structuring elements Y_a, Y_b and Y_c are still of the same size as the original element Y . But these atomic elements can be computed very efficiently, as all positions in the structuring element with an entry of $-\infty$ can be neglected during the computation of the morphological operation, because the generation of the minimum or the maximum would definitely eliminate the contribution of these elements. This results in a pure one-dimensional analysis of the signal in one of the three dimensions. This approach is similar to the tensor product decomposition applied in the field of linear filtering. In the case of tensor product wavelets, one also gets the anisotropy property, that results from the application of the structuring element Y .

If the size of the structuring element is again $2n + 1$ the computational complexity now results in $6n + 3$, since the application of a one-dimensional structuring element has to be performed three times. The elements Y_a, Y_b and Y_c could of course be decomposed again, but this is not reasonable in the one-dimensional case. The application of a one-dimensional structuring element with the size of $2n + 1$ results in $2n + 1$ operations. If this element would be decomposed into several smaller ones with an extent of three, the computational complexity would be $3n$ operations.

In the final result the opening of the data set is performed as follows:

$$X \circ Y = (((((X \ominus Y_a) \ominus Y_b) \ominus Y_c) \oplus Y_c) \oplus Y_b) \oplus Y_a \quad (14)$$

As these operations work on single lines, that are parallel to the x-, y-, or z-direction, a rearrangement of the data set is done between

most of the morphological operations. The idea is to have a set of vectors, that may be processed independently from each other. In the beginning of the process a bundle of vectors, that contain x-axis aligned data, are eroded and the new y-axis aligned vectors are constructed. These are eroded again and then transferred to z-axis aligned data. After an erosion of these data vectors, the inverse process is done with the dilation. This overall process is sketched in figure 6.

3.2 Parallelization

The overall structure can be easily parallelized. We have done this on a shared memory architecture using the POSIX thread standard. The first possibility for parallelization can be derived from Fig. 2. After reading the data set four threads are spawned to perform the opening operations for the four requested levels. For further computations these threads are joined afterwards. The following computational steps have not been parallelized, as the computational time required there is negligible in relation to the time needed for the opening operations. In consequence, the spawning of the threads probably would be more expensive than the gain in computation time.

These threads may again spawn some child threads. The erosion and dilatation operations, that are performed on a string can be parallelized by assigning a bundle of strings to each thread, which are then processed independently. In order to finish one of the morphological operations, all sub-threads have to be joined again. The number of sub-threads to be spawned can be specified by the user. If the analyzed volume has a size of 256^3 voxels, this would mean, that a maximum of $4 \cdot 256^2$ threads could be computing in parallel at one time.

Although this program has been implemented on a shared memory architecture it would also be suitable for a message passing mechanism, since the data locality of this problem is very high. Once the data set is copied, the different opening operators do not need to exchange data until they are finished and the operators working on the strings only need the data, that is coded into these strings.

4 Results

In order to show the usefulness of this method we have applied the morphological scale space analysis to four data sets. All data sets are CT scans. The first one is a scan of a cup, the second one represents a human head, the third one represents an engine block and finally the last one a simulated CT scan of a lung.

The cup data set (size $512 \cdot 512 \cdot 50$ voxels) is presented in order to illustrate the principal properties of the described algorithm. The main feature of the cup, that is used for the illustration of our algorithm, is the thickness of the wall. The thickness of the material increases from top to bottom and again becomes smaller in the very proximity of the bottom. In order to visualize this feature the necessary values have been adjusted as described in the first column of table 1.

The result of this analysis in the color plate. The images show the cup with all color components in image 7(g), the images 7(h) and 7(i) show the middle and the low frequency part of the volume separately. The image 7(h) with the green component illustrates the thickness of the cup walls. The color intensity is high at the bottom plate and at the upper parts of the walls, while the middle part is quite dark. This part of the cup is characterized by a thicker wall, as can be seen on image 7(i), which reflects the lowest frequency part. This image exhibits the most intense color values, where the green cup has its minimal values. In image 7(g), which contains all frequencies, one can clearly see this change in color. The noise

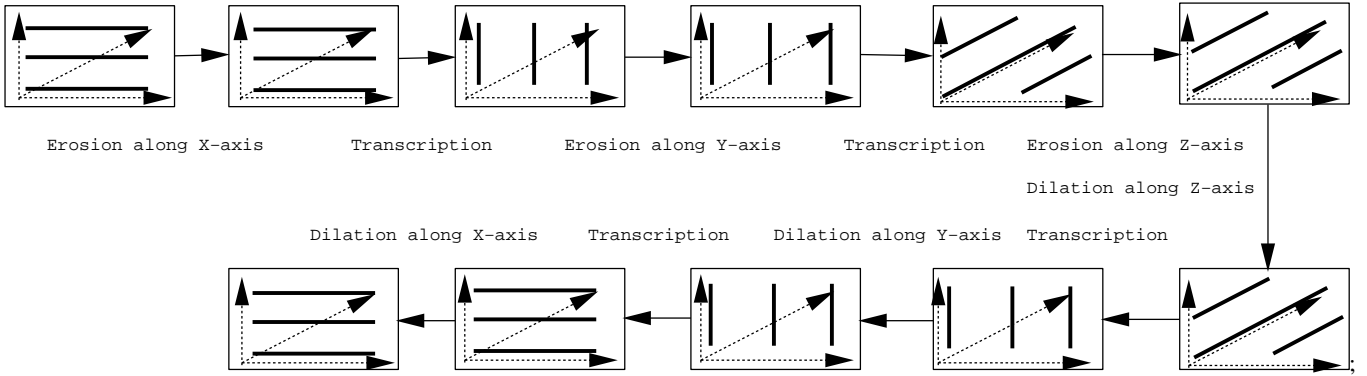


Figure 6: Data flow for one opening process

Parameter Name	Cup	Head	Engine	Lung
Element Size Level A	1	3	3	1
Element Size Level B	3	5	5	1
Element Size Level C	9	11	11	1
Element Size Level D	15	17	17	4
Red Threshold	0	30	0	0
Green Threshold	100	10	0	0
Blue Threshold	30	0	0	5
Red Scaling	0.7	1	3.5	10
Green Scaling	0.5	1.5	0.6	8
Blue Scaling	0.7	0.5	0.5	2
Red Absorption Influence	0.1	0.8	0.3	0.2
Green Absorption Influence	0.1	0.6	0.3	0.2
Blue Absorption Influence	0.1	0.3	0.05	0.1

Table 1: Parameters used for the data sets

Data Set	Resolution	Computational Time
Cup	512 · 512 · 50	158 sec
Head	256 · 256 · 113	80 sec
Lung	256 · 256 · 175	122 sec
Engine	256 · 256 · 110	81 sec

Table 2: CPU time for the analysis of the data sets

that surrounds the cup is part of the data set and caused by the CT scanner.

As a medical application we have investigated a CT scan of a head, that has a size of $256 \cdot 256 \cdot 113$ voxels. The morphological analysis has been applied here in order to increase the spatial impression of the volume rendering of the data set. The necessary parameters are listed in the second column of table 1. The head is shown in image 7(e). The plate behind the head is the pad, which supported the patient. The red part of the volume, that surrounds the head, again characterizes the low frequency structures of the volume. The green structures approximate the surface of the head in a more accurate manner. The anisotropy property again manifests itself at the low frequency structure, as the red region at the top of the head appears to be flattened. The combination of the different layers, that describe the head surface provide this enhanced spatial impression. In the front part of the head, that contains the mouth and the nose regions, many small details are reflected as high frequency blue structures, in contrast to the more homogeneous brain, that is colored in solid green. For comparison, a traditional volume rendering of the data set is displayed in image 7(f). This rendering provides more detailed information about the inner structure of the head, but the spatial impression of the overall head is quite low.

The third application of the morphological analysis is a scan of an engine block. Its consists of $256 \cdot 256 \cdot 110$ voxels and the applied parameters are listed in the third column of table 1. This engine block is characterized by various parts of different sizes. Our method automatically colors those parts of this engine block differently and thus reveals information about the component structure.

The image 7(a) shows the motor block rendered with all colors, while the images 7(b) and 7(c) separate the components of low frequency and middle frequency details. Two animations of the engine block are shown in the video. The first one rotates the engine block with all color components while the second one additionally cycles through the components to provide a better impression of the different detail levels.

The cylinders of the engine are colored in red as they are a very coarse structure in the engine block. The injection tubes consist of very thin metal parts, that are represented on the high frequency and fine granular structure level. This can be especially seen in image 7(b), where the injection tube is a very articulate feature, but it is nearly disappeared in the middle frequency scale in image 7(c). Most of the supporting structures of the motor block are represented at the middle level of detail, which is rendered in green. These structures can not be separated by a traditional volume rendering method, as they consist of nearly the same value. This type of data set seems to be the most suitable for our data analysis method.

The volume rendering of the simulated CT scan of the lung is shown in image 7(d). This data set has an extension of $256 \cdot 256 \cdot 175$ voxels. In this image the thickness of the airtubes is effectively visualized. What one can see from this artificial data set is that wall thickness of the main tube appears to be very small, where the walls of the first branch seems to be very thick. For the rendering of this image the prioritization order of the three frequency spectra has been inverted. In this case the lower frequencies suppress the higher ones.

All computations have been performed on a SGI Onyx with four R10000 processors running at 196 MHz. The required computational times for the data sets are listed in table 2. This time is the wall clock time that elapsed from the beginning of the data analysis until the RGBA values for the complete volume have been computed. The required CPU time increases linearly with the number of voxels to be processed, as could be expected from the complexity estimation. For standard data sets the required CPU time for the computation is below one and half minute and even for the cup, which is quite a large volumetric data set, it is below three minutes.

5 Conclusions

Constructing a scale-space method in analogy to the wavelet transformation using morphological operators has proven to be a valuable approach in volume visualization. The different results of the morphological analysis indicate boundary structures and they allow to distinguish objects of different characteristic sizes. In contrast to the application of traditional transfer-functions our method is easier to calibrate and it also incorporates spatial information. On the other hand, it is more application independent and more efficient than most of the special-purpose segmentation techniques known from pattern recognition. The separation of dimensions allows an efficient parallelization with a potentially fine granularity.


For future work we are thinking about incorporating aspects of region growing in order to generate surface models, that may be combined with the volume representation. The results of the morphological operators may also be used to provide the external forces for a deformable surface visualization technique. The different frequency levels found during the morphological analysis may guide the conversion process of the surface on different levels of resolution. There might also be the possibility to generate the optical properties of the volume by the morphological analysis in a more indirect way. This could be done by the selection of some predefined transferfunctions depending on the result of the morphological analysis.

Acknowledgements


Many thanks to Markus Blank of the institute for medical physics of the university Erlangen-Nürnberg for providing the artificial CT scan of the lung. Finally we would like to thank Gerold Herold, who provided the cup data set and Rüdiger Westermann, who generated the classical volume rendering image of the CT scan of the head.

References


- [1] J. Gauch and S. Pizer. Multiresolution Analysis of Ridges and Valleys in Grey-Scale Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):635–646, 1993.
- [2] K.-H. Hoehne and W. A. Hanson. Interactive 3D Segmentation of MRI and CT Volumes using Morphological Operations. *Journal of Computer Assisted Tomography*, 16(2):285–294, 1992.
- [3] P. Jackway and M. Deriche. Scale-Space Properties of the Multiscale Morphological Dilation-Erosion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1):38–51, 1996.
- [4] A. Kaufmann. *Introduction to Volume Visualization*. IEEE Computer Society Press, 1991.
- [5] M. Levoy. Display of Surfaces from Volume Data. *IEEE Computer Graphics and Applications*, pages 29–37, May 1988.
- [6] S. Mallat and W.L. Hwang. Singularity Detection and Processing with Wavelets. Technical report, Courant Institute of Mathematical Sciences New York University, march 1991.
- [7] S. Mallat and S. Zhong. Characterization of Signals from Multiscale Edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(7):710–732, Juli 1992.
- [8] S. Mallat and S. Zhong. Wavelet transform maxima and multiscale edges. In *Wavelets and their Applications*. Boston: Jones and Bartlett Publication, 1992.
- [9] S. Muraki. Multiscale 3D Edge Representation of Volume Data by a DOG wavelet. In *Proceedings Symposium on Volume-Visualization '94*, pages 35–42, 1994.
- [10] B. Nichols, D. Buttler, and J. Farrell. *Pthreads programming*. O'Reilly and Associates, 1996.
- [11] H. Niemann. *Pattern Analysis and Understanding*. Springer, 1990.
- [12] P. Saiviroonporn, A. Robotino, J. Zahajszky, R. Kikinis, and F. Jolesz. Real Time Interactive 3D-Segmentation. Technical report, Harvard Medical School and Brigham Women's Hospital, 1997.
- [13] Georgios Sakas and Stefan Walter. Extracting Surfaces from Fuzzy 3D-Ultrasound Data. In *ACM Computer Graphics (SIGGRAPH)*, pages 465–474, 1995.
- [14] G. Sapiro, R. Kimmel, D. Shaked, B. Kimia, and A. Bruckstein. Implementing Continuous-Scale Morphology via Curve Evolution. *Pattern Recognition*, 26(9):1363–1372, 1993.
- [15] S. Sternberg. Grayscale Morphology. *Computer Vision, Graphics and Image Processing*, pages 333–355, 1986.
- [16] E.J. Stollnitz, T. D. DeRose, and D. H. Salesin. *Wavelets for computer graphics*. Morgan Kaufmann, 1996.
- [17] R. van den Boomgard and A. Smeulders. The Morphological Structure of Images: The Differential Equations of Morphological Scale-Space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(11):1101–1113, November 1994.
- [18] R. Westermann. A Multiresolution Framework for Volume Rendering. In *Proceedings Symposium on Volume-Visualization '94*, pages 51–57, 1994.
- [19] R. Westermann and T. Ertl. A Multiscale Approach to Integrated Volume Segmentation and Rendering. In D. Fellenner and L. Szirmay-Kalos, editors, *Computer Graphics Forum (Eurographics'97 Conference Proceedings)*, number 3, pages 117–129. Eurographics, Blackwell, 1997.
- [20] X. Zhuang and R. Haralick. Morphological Structuring Element Decomposition. *Computer Vision, Graphics and Image Processing*, pages 370–382, 1986.




(a) Engine block, side view




(b) Engine block, high frequencies




(c) Engine block, middle frequencies




(d) Lung data set



(e) CT-head, morphological analysis




(f) CT-head, classical volume rendering



(g) Cup, color-coding of wall thickness



(h) Cup, middle frequencies



(i) Cup, low frequencies

Figure 7: Volume rendering applications of the multi-scale morphological analysis