



Centre for Image Analysis

Swedish University of Agricultural Sciences
Uppsala University

Licentiate Thesis No.8

Segmentation and Analysis of Volume Images, with Applications

Filip Malmberg

Licentiate Thesis

Centre for Image Analysis
Swedish University of Agricultural Sciences

2008



Contents

1	Introduction	4
2	Volume image acquisition	5
2.1	Computed tomography	5
2.2	Magnetic resonance imaging	6
3	Fundamental image analysis concepts	7
3.1	Digital images	7
3.2	Spatial filtering	8
3.2.1	Smoothing filters	9
3.2.2	Gradient extraction	9
3.3	Distance transforms	10
4	Image segmentation	11
4.1	Segmentation by thresholding	11
4.2	Watershed segmentation	12
4.3	Live-wire segmentation	12
4.4	Graph-cut segmentation	13
5	Applications	15
5.1	Analysis of volume images of fibrous materials	15
5.2	Interactive segmentation of medical volume images	16
6	Contributions	18
7	Discussion and suggestions for future work	21

Included papers

The following papers are included in this thesis:

- I. *Binarization of Phase Contrast Volume Images of Fibrous Materials*,
F. Malmberg, C. Östlund and G. Borgefors.
Submitted for journal publication.
- II. *Measuring Fibre-fibre Bonds in 3D images of Fibrous Materials*,
F. Malmberg, J. Lindblad, C. Östlund, K.M. Almgren and E.K. Gamstedt.
Accepted for publication in *Proceedings of 13th European Conference on Composite Materials (ECCM 13)*, Stockholm, Sweden, 2008.
- III. *A 3D live-wire segmentation method for volume images using haptic interaction*,
F. Malmberg, E. Vidholm and I. Nyström.
In *Proceedings of 13th International Conference on Discrete Geometry for Computer Imagery (DGCI 2006)*, Szeged, Hungary, 2006. LNCS vol. 4245, 2006, Pages 663 - 673.
- IV. *Minimal Cost-Path for Path-Based Distances*
R. Strand, F. Malmberg and S. Svensson.
In *Proceedings of 5th International Symposium on Image and Signal Processing and Analysis (ISPA 2007)*, Istanbul, Turkey, 2007. Pages 379-384.

In addition to the papers in this thesis, the author has also contributed to the following publication:

- I. *Comparison of stress-transfer mechanisms in paper sheets and composites made from resin-impregnated sheets*,
K.M. Almgren, E.K. Gamstedt, P. Nygård, F. Malmberg, J. Lindblad and M. Lindström.
In *Stress-Transfer Mechanisms in Wood-Fibre Composites*, Licentiate Thesis in Solid Mechanics, Royal Institute of Technology, Stockholm, Sweden, 2007.

The author has contributed significantly to the work performed in all the included papers. For Papers I-III, the author has had the major responsibility for method development, implementation, and writing. In Paper IV, the theoretical results were developed in close collaboration with Robin Strand and Stina Svensson, while the implementation and experiments were mainly performed by the author.

Chapter 1

Introduction

Digital image analysis is the field of extracting relevant information from digital images. Recent developments in imaging techniques have made 3-dimensional *volume images* more common. This has created a need to extend existing 2D image analysis tools to handle images of higher dimensions. Such extensions are usually not straightforward. In many cases, the theoretical and computational complexity of a problem increases dramatically when an extra dimension is added.

A fundamental problem in image analysis is image segmentation, i.e., identifying and separating relevant objects and structures in an image. Accurate segmentation is often required for further processing and analysis of the image can be applied. Despite years of active research, general image segmentation is still seen as an unsolved problem. This mainly due to the fact that it is hard to identify objects from image data only. Often, some high-level knowledge about the objects in the image is needed. This high-level knowledge may be provided in different ways. For fully automatic segmentation, the high-level knowledge must be incorporated in the segmentation algorithm itself. In interactive applications, a human user may provide high-level knowledge by guiding the segmentation process in various ways.

The aim of the work presented here is to develop segmentation and analysis tools for volume images. To limit the scope, the focus has been on two specific applications of volume image analysis: analysis of volume images of fibrous materials and interactive segmentation of medical images. The respective image analysis challenges of these two applications will be discussed. While the work has been focused on these two applications, many of the results presented here are applicable to other image analysis problems.

Chapter 2

Volume image acquisition

In this section, we briefly describe the imaging methods that were used to capture the volume image data used in this thesis. The author has not been involved in the image acquisition.

2.1 Computed tomography

Computed tomography (CT) is a non-invasive 3-dimensional imaging technique commonly used in medicine and material science. The technique is based on taking a large number of images of an X-ray beam transmitted through the sample from different directions. From these images, the internal 3D structure of the sample can be mathematically reconstructed.

The term *microtomography* is used for tomographic imaging with resolution in the micrometer range. To achieve this resolution, a highly parallel beam, preferably with high brilliance and penetration ability, is required. Synchrotron sources are capable of producing such X-rays. A slice from a microtomography image of a paper sample is shown in Figure 2.1.

According to [8], the most straightforward way of imaging with X-rays is to use beam *absorption*. This results in reconstructed images where the intensity of each image element corresponds to the density of the material present at that element. For weakly absorbing materials, it can be difficult to get enough contrast using beam absorption. Such materials may instead be imaged using *phase* contrast, which is more sensitive than absorption contrast. In phase contrast images, changes in refractive index between different materials are detected. In the reconstructed images, these changes are visible as bright and dark bands around object boundaries. Both absorption contrast and phase contrast effects may be present in an image. With X-ray microtomography, the balance between absorption contrast and phase contrast is determined by the distance between the sample and the imaging sensor.

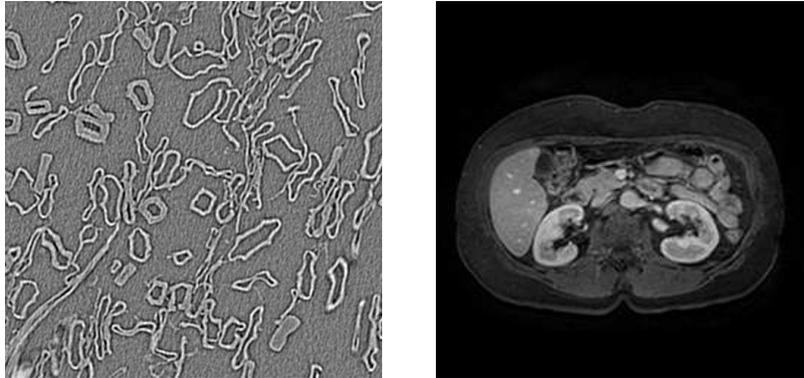


Figure 2.1: *Left: A slice from a X-ray microtomography volume image of a paper sample. Right: A slice from an MR volume image of the abdominal region of a human.*

2.2 Magnetic resonance imaging

Magnetic resonance imaging (MRI) is a widely used medical imaging technique. Compared to CT, MRI provides better soft-tissue contrast and has no known long-term health risks.

MRI imaging relies on the relaxation properties of excited hydrogen nuclei in water and lipids. The subject is placed in a strong static magnetic field, that causes the magnetic moment of the atom nuclei to align with the field direction. By applying a perturbation with another magnetic field in the form of a radio pulse, the spin of the nuclei is excited and causes the magnetic moments to point in the direction of the perturbing field. When the radio pulse is switched off, the spins will return to their equilibrium state. During this relaxation, radio frequency signals are emitted. These signals are detected by the MRI equipment, and can be used to calculate the time it takes for the magnetic moments to realign with the static magnetic field. Since the relaxation times vary between different tissues, contrast is obtained in the reconstructed images.

A slice from an MR image is shown in Figure 2.1.

Chapter 3

Fundamental image analysis concepts

In this section, digital images and the basic digital image analysis tools used in this thesis are presented.

3.1 Digital images

A gray-scale digital image can be described as a discrete, integer-valued, function of four variables, $f(x, y, z, t)$, where x , y , and z are spatial variables and t is time (e.g., in a movie sequence). The elements of the image are called *pixels* (picture elements) in 2D images, and *voxels* (volume elements) in volume images. See Figure 3.1. All images studied in this thesis are static, so the t variable will be disregarded.

A gray-scale image can be stored in a computer as an array of integers. Typically, 8 bits are used to represent each integer, which means that $2^8 = 256$

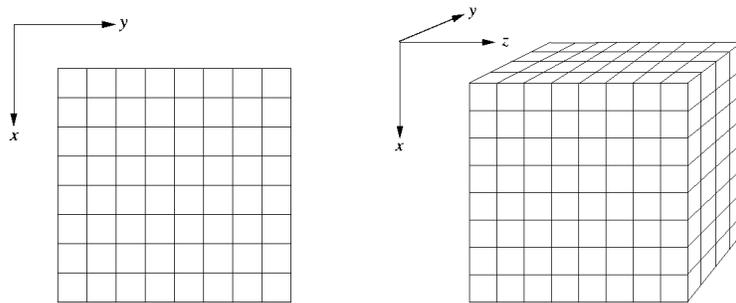


Figure 3.1: *Left: A 2-dimensional image consisting of square pixels. Right: A volume image consisting of cubic voxels.*



Figure 3.2: *Adjacency relations for pixels. (a) Edge neighbors. (b) Vertex neighbors.*

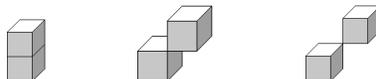


Figure 3.3: *Adjacency relations for voxels. (a) Face neighbors. (b) Edge neighbors. (c) Vertex neighbors.*

different gray-levels can be represented. However, any number of bits may be used. In a *binary* image, only 1 bit is used to represent each element. This means that only two different levels, black or white, can be represented.

Adjacency relations between image elements are important for many image analysis algorithms. Adjacency can be defined in a number of different ways.

A pixel in a 2D image has four horizontal and vertical neighbors, called *edge neighbors*, and four diagonal neighbors called *vertex neighbors* (see Figure 3.2). The set of all edge neighbors of a pixel is called the *4-neighborhood*, and the union of this set and all the vertex neighbors is called the *8-neighborhood*.

A voxel in a 3D image has three different kinds of neighbors: *face neighbors*, *edge neighbors*, and *vertex neighbors* (see Figure 3.3). As in the 2D case, we combine these types of neighbors to define different neighborhoods. Common neighborhood definitions are the *6-neighborhood* that consists of the face neighbors, and the *26-neighborhood* that consists of all face, edge, and vertex neighbors.

3.2 Spatial filtering

Spatial image filtering may be used for, e.g., noise reduction and edge enhancement. A filter mask, with weights defined over a neighborhood \mathcal{N} , is moved over all elements in the image. At each element, the values of the neighbors and the weights of the filter mask are used to compute an output value, the filter *response*. For *linear* filters, the response is a sum of products of the filter weights and the values of the corresponding image elements:

$$g(\mathbf{x}) = \sum_{\mathbf{s} \in \mathcal{N}} w(\mathbf{s})f(\mathbf{x} + \mathbf{s})$$

where w is the filter mask, f is the input image and g is the output image. Other combinations of the mask weights and the element values may also be used, in which case the filter is called *non-linear*.

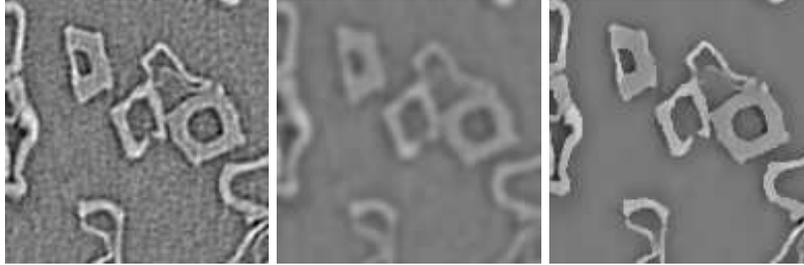


Figure 3.4: *Using smoothing filters to suppress image noise. Left: Original noisy image. Middle: A Gaussian smoothing filter reduces noise, but also blurs edges in the image. Right: SUSAN-filtering reduces noise while maintaining edges.*

3.2.1 Smoothing filters

Smoothing filters may be used to suppress high frequency noise in images. A commonly used smoothing filter is the Gaussian filter, where the filter mask is a discrete approximation of a Gaussian function with standard deviation σ :

$$w(\mathbf{s}) = e^{-\frac{1}{2}\|\mathbf{s}\|^2/\sigma^2}$$

To obtain a good approximation of the Gaussian function, the size of the filter mask should increase with the value of σ .

A problem with using smoothing filters for noise reduction is that edges in the image are also blurred, and thus important information may be lost. Edge preserving smoothing algorithms attempt to reduce image noise while preserving strong edges. One such approach is SUSAN-filtering [9] where two Gaussian functions are used, one for the spatial domain and one for intensity:

$$g(\mathbf{x}) = \frac{\sum_{\mathbf{s} \in \mathcal{N}, \mathbf{s} \neq \mathbf{x}} f(\mathbf{x} + \mathbf{s}) e^{-\left(\frac{\|\mathbf{x}-\mathbf{s}\|^2}{2\sigma^2} + \frac{(f(\mathbf{x}-\mathbf{s}) - f(\mathbf{x}))^2}{t^2}\right)}}{\sum_{\mathbf{s} \in \mathcal{N}, \mathbf{s} \neq \mathbf{x}} e^{-\left(\frac{\|\mathbf{x}-\mathbf{s}\|^2}{2\sigma^2} + \frac{(f(\mathbf{x}-\mathbf{s}) - f(\mathbf{x}))^2}{t^2}\right)}}$$

In the above formula, σ is the standard deviation in the spatial domain and t is the standard deviation in the intensity domain. The weights of the filter are thus attenuated, depending on the magnitude of the local gray-scale differences. A comparison between Gaussian smoothing and SUSAN-filtering is shown in Figure 3.4.

3.2.2 Gradient extraction

Edge information is often used in image segmentation to identify region boundaries. A commonly used edge indicator is the gradient magnitude $\|\nabla f\|$ of the image. The gradient of an image can be approximated using, e.g., centered differences:

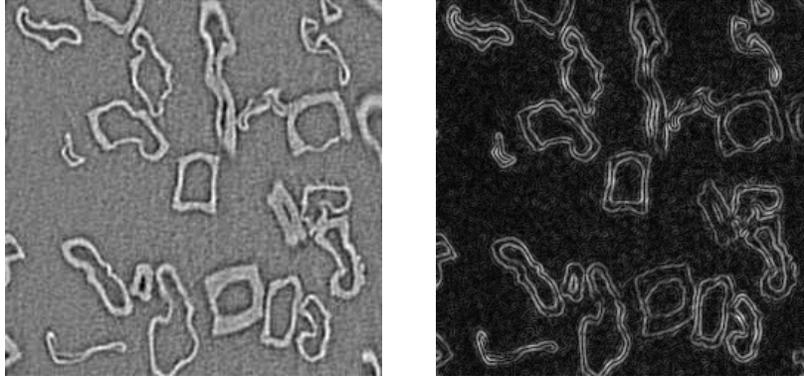


Figure 3.5: *Left: Original image. Right: Gradient magnitude.*

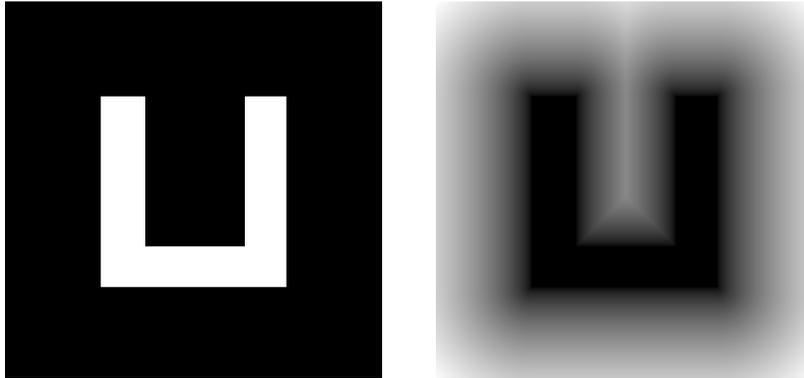


Figure 3.6: *Left: A binary image. Right: Corresponding distance transform.*

$$\nabla f \approx \begin{pmatrix} \frac{f(x+1,y,z) - f(x-1,y,z)}{\Delta x} \\ \frac{f(x,y+1,z) - f(x,y-1,z)}{\Delta y} \\ \frac{f(x,y,z+1) - f(x,y,z-1)}{\Delta z} \end{pmatrix}$$

The effect of the gradient magnitude operator is illustrated in Figure 3.5.

3.3 Distance transforms

Distance transforms (DTs) are useful in many applications of image analysis. A DT is applied to a binary image and computes, for each background element, the distance to the closest object element. The result of the DT is a *distance map*, a gray-level image containing the computed distance values. See Figure 3.6.

Chapter 4

Image segmentation

One of the most challenging problems in image analysis is segmentation, i.e., partitioning an image into objects of interest and background. Segmentation is often needed before further analysis of an image can be performed.

Fully automatic segmentation is desirable in most applications, but is not always achievable. In applications where accurate segmentation is critical, such as in medical image analysis, unsupervised automatic segmentation methods may not be robust enough. Manual segmentation of large volume images, on the other hand, is very time consuming and error-prone.

According to [7], the segmentation task can be divided into two steps: *recognition* and *delineation*. Recognition is the task of roughly determining where in the image the objects are located, while delineation consists of determining the exact extent of the object. Human users outperform computers in most recognition tasks, while computers are often better at delineation. Semi-automatic segmentation methods take advantage of this fact by letting a human user guide the recognition, while the computer performs the delineation. A successful semi-automatic method minimizes user interaction time, while maintaining tight user control to guarantee the quality of the result.

The best approach to image segmentation may vary between different applications. The choice between manual, semi-automatic or fully automatic methods depends on, e.g., the quality of the images, the number of objects that need to be segmented, the amount of available user time, and the required accuracy of the segmentation.

In this section, the different image segmentation methods used in this thesis are presented.

4.1 Segmentation by thresholding

One of the most straightforward methods for image segmentation is *gray-level thresholding*. Thresholding classifies each image element as object or background, depending on whether its intensity value is above or below some thresh-

old value t . This simple method may be sufficient for images where the gray-levels of the object and the background are clearly different. Even in simple images, however, the gray-level distributions of the object and the background often overlap due to shading variations and noise.

Hysteresis thresholding [3] is a slightly more sophisticated segmentation method. In this method, two threshold values t_{high} and t_{low} are chosen. Elements whose gray-level is above t_{high} are labeled as object, and elements with gray-levels below t_{low} are labeled as background. Elements with gray-level in the range $[t_{low}, t_{high}]$ are labeled as unknown. The unknown regions are then resolved in the following way: If a region of connected unknown elements touches at least one foreground element, the region is labeled as object. Otherwise it is labeled as background.

4.2 Watershed segmentation

Watershed segmentation [11] is a method for partitioning a gray-scale image into regions, based on the topology of the image.

The concept of watershed segmentation is most easily understood by interpreting the gray-level image as a landscape, where bright regions correspond to hills, and dark regions correspond to valleys. If we now let it rain over this landscape, the valleys will start to get filled by water. When the water reaches a level where two different basins are about to meet, a watershed is built. This prevents water collections belonging to different local valleys from merging. When the whole image landscape has been submerged, all image elements will be assigned to a basin of water originating from a local valley. Although the analogy above uses a 2D image, the watershed algorithm may be applied to images of any dimension.

The watershed algorithm can be applied to different types of image information, such as gray level, gradient magnitude or a distance transform, to achieve different results.

Since the watershed algorithm detects all local minima in the image intensity landscape, it often results in *over-segmentation*. This may be reduced by smoothing the image prior to watershed segmentation to remove some local minima, or by merging “similar” regions in a post-processing step.

4.3 Live-wire segmentation

Live-wire [7, 6] is a semi-automatic segmentation method for 2D images. The basic idea of the live-wire method is that the user places a seed-point on the boundary of an object. As the user moves the cursor in the image, a path (live-wire) is displayed in real-time from the current position of the cursor back to the seed-point. The wire is attracted to edges in the image and it will snap onto the boundary of the object. If the user moves too far from the original seed-point, the wire might snap onto an edge that does not correspond to the desired

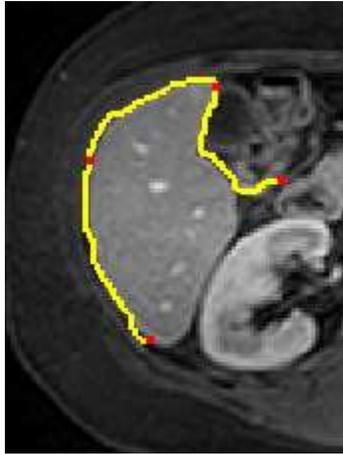


Figure 4.1: *Liver segmentation in an MR image with Live-wire. The user interactively places seed-points along the liver boundary. As the user moves the cursor, the shortest path from the last seed-point to the current cursor position is displayed in real-time.*

object. When this happens the user can move the cursor back a little and place a new seed-point. The wire from the old seed-point to the new then freezes, and the tracking continues from the new seed-point. In this way, the entire object boundary can be traced with a rather small number of live-wire segments. See Figure 4.1. The live-wire algorithm is based on shortest-path calculations in a graph-representation of the image. For every edge in the graph, a cost is assigned to represent the likelihood that the edge belongs to a desired boundary in the image. When the user places a seed-point, Dijkstra's algorithm [12] is used to compute the optimal path to the seed-point from all other points in the image. Once these paths have been computed, it is possible to display the live-wire in real-time with virtually no computational cost as the user moves the cursor through the image.

With a slight modification to Dijkstra's algorithm, more than one seed-point may be used to compute the shortest path. In [5] this modified algorithm is referred to as the *image foresting transform*. In Paper III, the image foresting transform is used to efficiently compute multiple live-wires between two user-defined discrete curves.

4.4 Graph-cut segmentation

Graph cut segmentation [2] is an image segmentation method based on combinatorial optimization techniques. The method is applicable to images of any dimension and gives a binary partitioning of the image into background and object.

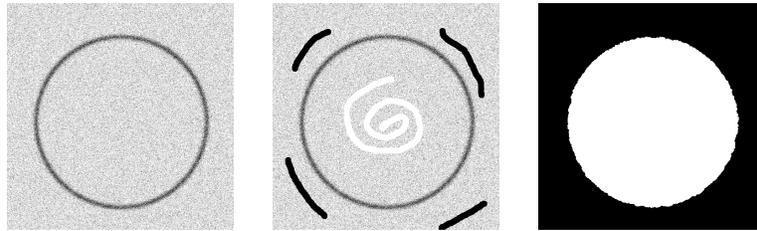


Figure 4.2: *Graph cut segmentation of a synthetic image. Left: Original image. Middle: User initialization. Right: Segmentation result.*

In the graph-cut method, the image is represented as a graph where image elements correspond to nodes and paths between adjacent elements correspond to graph edges. Each graph edge is assigned a non-negative cost, that represents the likelihood that the edge belongs to the boundary of the object of interest. Two special nodes are added to the graph, the *source* node and the *sink* node. Image elements that are a priori known to belong to the object are connected to the source node with zero cost edges. Similarly, elements that are known to belong to the background are connected to the sink node. Given this input, the graph-cut method computes a *cut* on the graph that separates the source node from the sink node, thus associating each image element with either the object or the background. The computed cut is optimal, in the sense that the sum of the cost of all edges that cross the boundary between the object and the background is minimal.

Graph cuts are commonly used in interactive segmentation, in which case the user initializes the graph by manually labeling some elements in the image as either object or background. As illustrated in Figure 4.2, relatively sparse input is often sufficient. In some applications, sufficiently good initial labels may also be found by an automatic segmentation algorithm.

Chapter 5

Applications

In this section we describe the applications of volume image analysis that the work presented in this thesis has contributed to. The respective image analysis challenges of both applications are discussed.

5.1 Analysis of volume images of fibrous materials

Wood fibres are used in many types of materials. The most common such materials are paper and board, which consist of a dense network of pulp fibres. Another, quite new application for wood fibres is composite materials, where a network of fibres is used to reinforce a polymer matrix.

In this project, we develop image analysis tools for studying the 3-dimensional micro-structure of fibrous materials. On a microscopic scale, these materials consist of wood fibres that bind together in a complex network. The properties of this network are related to the macroscopic properties of the material, so measurements of the micro-structure can be used to increase our understanding of the material and how it can be improved.

Recently, X-ray microtomography has been successfully used to capture high resolution volume images of fibrous materials [1, 8]. Automated analysis of such images can give a lot of useful information about the properties of the material. A surface rendering of a volume image of a pulp-fibre composite material, captured with X-ray microtomography, is shown in Figure 5.1.

The analysis of such images presents a number of challenges. The large number of fibres in the images makes manual or semi-automatic segmentation of the fibres unpractical. Methods for segmentation and analysis of the images should therefore be as automated as possible. The large amounts of data makes computational efficiency of the developed methods an issue.

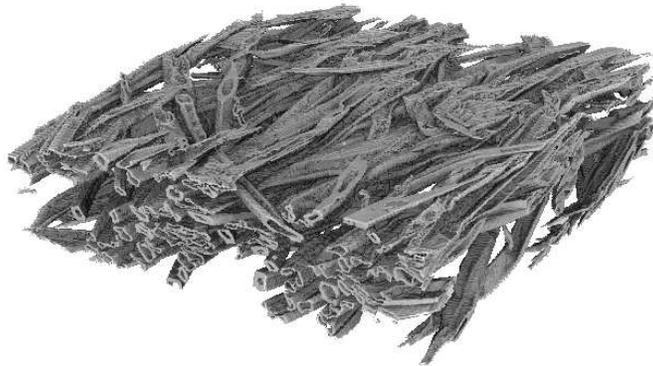


Figure 5.1: *A surface rendering of fibres in a pulp-fibre composite material, captured with μ CT. The polymer matrix surrounding the fibres is not shown.*

5.2 Interactive segmentation of medical volume images

Segmentation of medical images requires high accuracy. In many cases, however, only a small number of individual objects need to be segmented. This makes interactive semi-automatic segmentation methods attractive for this application.

Many semi-automatic segmentation methods exist for 2D images. However, extending these to 3D is often a non-trivial task. In addition to the problem of extending the segmentation algorithm itself to higher dimensions, the way the user interacts with the data may also need to be modified. Interaction with volume images is difficult, partly because most input devices, such as a computer mouse, are designed for 2D interaction only. To facilitate efficient interaction with volume data, we use an input device that supports 3D input and haptic feedback. This gives new possibilities to create interfaces where the user can simultaneously explore and manipulate volume images in an intuitive way. Our aim in this project is to investigate how these new possibilities can be used to extend existing semi-automatic segmentation methods to also handle volume images efficiently. The input device, together with a stereoscopic display, is shown in Figure 5.2.



Figure 5.2: *A workstation for interactive segmentation of volume images, featuring a 3D input device with haptic feedback and a stereoscopic display.*

Chapter 6

Contributions

In this section, the methods and results for the appended papers are presented briefly.

Paper I: Binarization of Phase Contrast Volume Images of Fibrous Materials

Many image analysis methods require a binary image as input. Segmenting volume images of fibrous materials into fibre and background is therefore an important pre-processing step. We call this step *binarization*.

In absorption mode images, the intensity of each image element corresponds to the density of the material present at that element. Such images can therefore, after some pre-processing, be binarized using intensity thresholding.

Phase contrast images, on the other hand, are conceptually harder to binarize than absorption images, since most individual image elements contain no information about which material they belong to.

In Paper I, we present a method for binarizing phase contrast volume images of fibrous materials. The proposed method first identifies the interface bands at the boundaries between fibre and background using thresholding. Graph-cut segmentation is then used to obtain a final binarization from the segmented interface bands. The method is shown to produce better results than a standard method based on edge preserving smoothing and hysteresis thresholding.

Paper II: Measuring Fibre-fibre Contact in 3D Images of Fibrous Materials

The degree of fibre-fibre bonding plays a fundamental role in the mechanical properties of paper materials and fibre mats. A common measure is the relative bonded area, RBA, defined as the ratio of bonded area to the total surface area of the fibres. It is possible that the stress transferring fibre-fibre bonds could be

of importance for mechanical properties also for pulp-fibre composites. Existing methods for measuring fibre-fibre bonding are, however, not applicable to pulp-fibre composites, due to the properties of the polymeric matrix that embeds the fibres in such materials.

In Paper II, an automatic method for measuring fibre-fibre contact in binary volume images of fibrous materials is presented. In contrast to previous methods for measuring RBA, the proposed image analysis method is applicable to both paper and composite materials. The method is tested on X-ray microtomography volume images of pulp-fibre composites and is shown to successfully detect differences in the amount of fibre-fibre contact between the samples.

Paper III: A 3D Live-wire Segmentation Method for Volume Images Using Haptic Interaction

In Paper III, we present an extension of the live-wire method to 3D images. In the 2D live-wire method, the user places seed-points that are connected by minimum cost-paths. Our idea for a 3D extension is to let the user draw 3D live-wire curves which are connected by discrete surfaces that follow the boundaries of the underlying object, a process we call *bridging*. The aim is to segment entire objects by drawing a relatively small number of live-wire curves on the boundary of the object. The proposed bridging algorithm uses the image foresting transform to create a large number of live-wire curves between two user defined curves. The live-wire curves are then connected by polygons, which are rasterized to obtain a tunnel-free discrete surface that closely matches the underlying objects in the volume image.

Since the publication of Paper III, we have become aware of graph-cut segmentation, a method that elegantly solves the problem of finding minimum-cost discrete surfaces in images of any dimension. The graph-cut method has many nice features, and in many cases it may be a better solution than the bridging algorithm presented in Paper III.

Paper IV: Minimal Cost-Path for Path-Based Distances

The Euclidean distance function is used in many image analysis applications since it has minimal rotational dependence. However, in some applications, path-based distances, such as the city-block and chessboard distances, are still preferable. One such example is the computation of the constrained minimal-cost path between two points in an image. Finding constrained minimal-cost paths is an important problem in image analysis, with applications to e.g. image segmentation [7]. For path-based distance functions the problem can be solved efficiently in any number of dimensions using e.g. Dijkstra's algorithm [4]. However, for path-based distances, the solution is not necessarily unique.

To decrease the rotational dependency of path-based distances, different weights can be used for different neighbors. Recently, neighborhood sequence distances, where the neighborhood relation is allowed to change between different steps in the path, have been studied [10].

In Paper IV, we study path-based distances from the perspective of generating unambiguous shortest-paths. We establish a relation between the shape of the ball generated by a distance function and the size of the set of minimal-cost paths. Our results suggest that the size of the set of minimal-cost paths decreases as we minimize the rotational dependency of the distance function.

The results in Paper IV are directly applicable to live-wire segmentation. By optimizing the distance function used to compute the live-wire curves according to the criteria derived in Paper IV, more regular live-wire curves can be obtained in image regions with weak edges.

Chapter 7

Discussion and suggestions for future work

Automatic analysis of fibre images is a promising tool. Microtomography imaging techniques are continually improving, and desktop 3D scanners are now approaching the high resolutions needed for analysis of individual fibres. The ability to capture high resolution images of fibrous materials without the need for a synchrotron source will greatly increase the availability of such images, and therefore improve the practical use of the image analysis tools presented here.

Many measurements on the fibre network would be much easier to compute if we could segment all individual fibres in the images. This is, however, a very difficult task. This is due both to the fact that the fibre network can be very dense and complicated, and that the fibres themselves can be so damaged and deformed during the manufacturing process that it is hard to find common characteristics for the shape of a fibre. In the work presented here, we have instead opted for indirect measurement methods, that do not require segmentation of the individual fibres. For future work, however, robust segmentation of individual fibres is a high priority.

In the field of interactive segmentation, our future work will focus on graph-cut segmentation. There are several interesting topics to be explored. First, we plan to investigate how 3D input and haptic feedback can be used to create an efficient user interface for initializing graph-cut segmentation. We would also like to explore various ways of reducing the computation time of the method, to achieve better interactive feedback for segmenting volume images.

Bibliography

- [1] C. Antoine, P. Nygård, Ø Weiby Gregersen, R. Holmstad, T. Weitkamp, and C. Rau. 3d images of paper obtained by phase-contrast x-ray microtomography: image quality and binarization. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 490(1-2):392–402, 2002.
- [2] Y. Boykov. Graph cuts and efficient n-d image segmentation. *International Journal of Computer Vision*, 70(2):109–131, 2006.
- [3] J. F. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, 1986.
- [4] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [5] A. X. Falcão and F. P. G. Bergo. Interactive volume segmentation with differential image foresting transforms. *IEEE Transactions on Medical Imaging*, 23(9), September 2004.
- [6] A. X. Falcão, J. K. Udupa, and F. K. Miyazawa. An ultra-fast user-steered image segmentation paradigm: Live wire on the fly. *IEEE Transactions on Medical Imaging*, 19(1):55–62, January 2000.
- [7] A. X. Falcão, J. K. Udupa, S. Samarasekera, and S. Sharma. User-steered image segmentation paradigms: Live wire and live lane. *Graphical Models and Image Processing*, (60):223–260, 1998.
- [8] E.J. Samuelsen, O.W. Gregersen, P.J. Houen, T. Helle, C. Raven, and A. Snigirev. Three-dimensional imaging of paper by use of synchrotron x-ray microtomography. *Journal of Pulp and Paper Science*, 27(2):50–53, 2001.
- [9] S.M. Smith and J.M. Brady. SUSAN- a new approach to low level image processing. *Int. Journal of Computer Vision*, 23(1):45–78, 1997.
- [10] R. Strand. Weighted distances based on neighbourhood sequences. *Pattern Recognition Letters*, 28(15):2029 – 2036, 2007.

- [11] L. Vincent and P. Soile. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, 1991.

Acknowledgments

I would like to thank the following people who, in various ways, contributed to this thesis:

- My supervisors Gunilla Borgefors, Joakim Lindblad and Catherine Östlund for their scientific guidance and support, and for giving me the chance to do research in a very interesting field.
- Ingela Nyström and Stina Svensson for valuable support and good collaborations.
- My friends and colleagues at CBA, for making it such a nice place to work at.
- My family: Annika, Anders, Anna, Martin and Jocke, for constant love and support.