

On 3D Fibre Measurements of Digitized Paper

from Microscopy to Fibre Network

Mattias Aronsson

*Centre for Image Analysis
Uppsala*

**Doctoral Thesis
Swedish University of Agricultural Sciences
Uppsala 2002**

Acta Universitatis Agriculturae Sueciae
Silvestria 254

ISSN 1401-6230
ISBN 91-576-6338-6
© 2002 Mattias Aronsson, Uppsala
Printed at the Uppsala University Printing and Media Office 2002

Abstract

Mattias Aronsson. *On 3D Fibre Measurements of Digitized Paper*. Doctoral Thesis
ISSN 1401-6230, ISBN 91-576-6338-6

Paper is a material that suits many different purposes, and is therefore used extensively. To be able to optimize the manufacturing parameters during the production in a paper mill, it is of great importance to know as much as possible about the material. Paper exists in many different qualities, and different paper grades can have very different properties. There already exists many techniques to measure paper properties, but nearly all are based on computerized image analysis of two spatial dimensions, or not based on image analysis at all. Recently, the processing power of computers has reached a level which makes it possible to analyze huge data sets, by 3D computerized image analysis. Most grades of paper consists of fibres connected in a 3D network, together with fillers and surface coatings. We have focused on the 3D network of fibres, to provide techniques to look into, and perform calculations on this network. Advantages of using image analysis is that if we can transform samples of paper into an accurate digital format, a computer can then be programmed to analyze some of the properties not easily accessible by other means. Both mechanical strength and opacity depend on the geometry of the fibre network, so the digital model could be useful for further analysis of these properties. Disadvantages are that the conversion from physical to digital form will introduce artefacts and this conversion is often time-consuming. Our work has been to develop a method to construct 3D digital models of paper and to develop some useful measurements. The digital models consist of efficient representations that enable our measurements and also provide a good base for further development. The measurements consist of a detailed fibre analysis, a few fibre network measures, and experiments to measure the fibre pore network. Both 2D and 3D based fibre segmentation algorithms have been implemented. The image capturing process is crucial, but we still believe that the core problem is to improve the fibre segmentation. Especially, a good reconstruction of the network requires a high percentage of the fibres to be found, something that need further work to be accomplished. This thesis will emphasize the image analysis part of the problem, even if there are as much to be said about paper as a material.

Key words: volume images, image registration, filter design, fibre network, individual fibre segmentation, fibre properties estimation, visualization

Author's address: Centre for Image Analysis, Lägerhyddsvägen 17, 752 37 Uppsala, Sweden



CALVIN AND HOBBS © Watterson. Reprinted with permission of
UNIVERSAL PRESS SYNDICATE. All rights reserved.

L^AT_EX

Contents

On 3D Fibre Measurements of Digitized Paper

1	Introduction	9
1.1	Papermaking	9
1.2	The “3D tracking of fibres in paper” project	11
1.3	Paper research	12
1.4	Why 3D image analysis?	15
1.5	Digital 3D model of paper	17
2	Data acquisition	19
2.1	Sample preparation and scanning	19
2.2	Light Microscopy	19
2.3	Scanning Electron Microscopy	20
2.4	X-ray micro tomography	22
2.5	Combining micro tomography and SEM	22
2.6	Imaging conclusions	24
3	Methods	26
3.1	Implementing efficient 3D image analysis algorithms	26
3.2	Raw data	29
3.3	Registering the set of images into a 3D image	29
3.4	Filtering and pre-processing	33
3.5	Resolution enhancements	34
3.6	Where is the surface?	34
3.7	3D pore distribution	36
3.8	Segmenting individual fibres	37
3.9	Measurements	47
3.10	Rendering 3D images	49
4	Summary of the included papers	56
4.1	Paper I: Comparison of two different approaches for paper volume assembly	56
4.2	Paper II: Minimizing scanning electron microscope artefacts by filter design	56
4.3	Paper III: Slice based digital volume assembly of a small paper sample	57
4.4	Paper IV: 2D Segmentation and labelling of clustered ring-shaped objects	57
4.5	Paper V: Ring-shaped object detector for non-isotropic 2D images using optimized distance transform weights	58
4.6	Paper VI: Some measurements of fibres in volume images of paper using medial representations detected on the distance transform	58
4.7	Paper VII: Estimating fibre twist and aspect ratios in 3D voxel volumes	59
4.8	Paper VIII: Using distance transform based algorithms for extracting measures of the fibre network in volume images of paper	59
4.9	Paper IX: 3D Pore distribution estimation in digitized paper samples	60

5	Conclusions and further results	61
6	Other publications and conferences	70
7	Acknowledgements	71

Papers I–IX

Papers appended to the thesis

The thesis is based on the following articles. Paper I was presented at the *Swedish Society for Automated Image Analysis (SSAB)* in Halmstad Sweden 2000, Paper II was published in the *Journal of Microscopy (JMS)* 2001, Paper III was presented at the *Scandinavian Conference on Image Analysis (SCIA)* in Bergen Norway 2001 as a poster, Paper IV was presented at the *International Conference on Image Processing (ICIP)* Rochester USA 2002 as a poster, Paper V was published in the *Nordic Pulp and Paper Research Journal (NPPRJ)* in 2002, Paper VI was presented at the *Computer Vision and Pattern Recognition (CVPR)* conference in Hawaii Kauai USA 2001 as a technical sketch, a revised version of Paper VII has recently been submitted to the *IEEE Transactions on System Man and Cybernetics*, intended for a special issue about 3D image analysis. Paper VIII was presented at the *International Conference on Pattern Recognition (ICPR)* Québec City, Canada 2002 as a poster. Paper IX is included as a manuscript. Paper I and IX were not reviewed, and we are still waiting for a definite answer for Paper VI, while the other six papers have been evaluated and accepted by international reviewers. The degree of contribution for each of the included papers for the author of this thesis is approximately: Paper I – 50%, Paper II – 70%, Paper III – 80%, Paper IV – 80%, Paper V – 70%, Paper VI – 40%, Paper VII – 40%, Paper VIII – 100%, and Paper IX – 80%. As many of the papers are now copyright material owned by the publisher, we have asked for permission to include these in this thesis, and this was granted.

- I Comparison of two different approaches for paper volume assembly; *Aronsson, M. and Fayyazi, A.* (Proceedings)
- II Minimizing scanning electron microscope artefacts by filter design; *Aronsson, M., Sävborg, Ö., and Borgefors, G.* (J. Microscopy 2001)
- III 2D Segmentation and labelling of clustered ring shaped objects; *Aronsson, M. and Borgefors, G.* (Proceedings)
- IV Ring-shaped object detector for non-isotropic 2D images using optimized distance transform weights; *Aronsson, M. and Sintorn, I.-M.* (Proceedings)
- V Slice based digital volume assembly of a small paper sample; *Aronsson, M., Henningsson, O., and Sävborg, Ö.* (NPPRJ 2002)

- VI Some measurements of fibres in volume images of paper using medial representations detected on the distance transform; *Svensson, S. and Aronsson, M.* (Proceedings, Technical Sketch)
- VII Using distance transform based algorithms for extracting measures of the fibre network in volume images of paper; *Svensson, S. and Aronsson, M.* (Submitted paper)
- VIII Estimating fibre twist and aspect ratios in 3D voxel volumes; *Aronsson, M.* (Proceedings)
- IX 3D Pore distribution estimation in digitized paper samples; *Aronsson, M. and Borgefors, G.* (Manuscript)

1 Introduction

This thesis will tell you about the efforts made to find and measure fibres in paper samples, by using computerized image analysis. There are numerous steps involved to implement these measurements, moreover we need to convert the paper sample into a digital format suitable for computer based analysis. Thus a lot of programming is needed, and especially the knowledge how to implement image analysis methods on a computer in a robust and efficient way is essential. The focus is on the image analysis, but some basic paper and fibre knowledge is necessary to understand some of the properties of the image material we use. As you will see, we have solved parts of the problem, but there are more to be done before this application can be packaged as a “Plug and Play” piece of software.

Basically any problem suitable for image analysis involves the following steps: (1) Decide what imaging device is suitable and make the arrangements to build a working setup. This step transforms a continuous image from a sensor into a digital image that we can manipulate with a computer. (2) Pre-processing of images. This may be noise attenuation, lens corrections of the optical system, attenuating artefacts from the imaging, or other adjustments that make the following steps more robust. (3) Segmentation of the digital image, i.e., find the objects of interest. (4) Represent the objects found in a more convenient format, to facilitate calculations of properties/features. (5) Measure properties of interest. (6) Interpretation of results.

Some of the challenges for this specific project have been the handling of large amounts of data and the segmentation of individual fibres, which is often the most difficult part of the problem. It is no coincidence, as the segmentation step basically gives the computer the ability to see and interpret the image, which is no small feat. As always with computers and programming, there are lots of special details that need proper care to make the source code behave correctly.

1.1 Papermaking

The paper sample we have used for most of our analysis was made in the paper machine KM7 of the Skoghall Mill outside Karlstad, Sweden. Schematics over the machine can be seen in Figure 1. The length of this machine is around 250 metres, and the new and bigger KM8 is almost 300 metres long! Also the width of the KM8 paper reels are very wide, 8.1 metres. You may realize the importance to minimize the number of *paper jams*, since when this occurs, it takes a lot of time to reset the process and attain full speed again. Another somewhat amazing detail about paper manufacturing is the speed of paper formation. Modern, large paper machines manufacturing lightweight grades, e.g., newsprint, can run at speeds up to 40 m/s (i.e. 144 kph or 90 mph!). Heavier paper grades like board can not be manufactured at such speeds, KM7 in Skoghall runs at about 9 m/s and the new and bigger KM8 at about 13 m/min. During a “flawless” year KM8 will produce more than 300,000,000 kg of paper, or stated differently, produce more than 10 kg of paper in a single second.

The environmental aspects of the paper manufacturing process have gained more

and more interest in recent decades, so now paper material is extensively recycled and reused when making new paper. The chemicals needed are also recycled internally during the process, minimizing the waste. As background information, it may be of interest to know the basics in paper manufacturing. Therefore we will give a short introduction to this subject. When the trees are harvested, the branches are removed and then the tree trunks are debarked and chopped into small wood chips. There are basically three different processes to convert the wood chips into free fibres: (a) Mechanical pulping uses direct grinding of the wood logs between rotating cylinders to get free fibres; (b) Thermo-mechanical pulping (TMP) is a process in which chips are mechanically refined after being heated in a pressurized environment; (c) Chemical pulping, where the wood chips are delignified by chemical solvents in a digester.

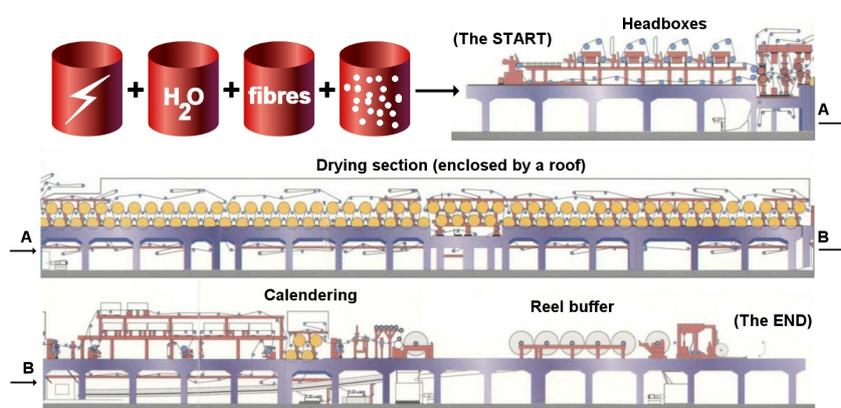


Figure 1: The Skoghall Mill KM7 paper machine close to Karlstad, Sweden.

As the fibres and some of the other material still left in the diluted fibre suspension may have large colour variations, a bleaching step may be used to improve the colour consistency. Common chemicals used are chlorine dioxide, hydrogen peroxide, ozone, or oxygen. The paper forming takes place in the paper machine and in the first step at the headbox, the much diluted suspension of water and fibres (less than 1% fibre material) is injected either in-between an upper and lower moving wire, or onto a moving wire. The “trick” of using different speeds of the wire and the injected fluid greatly influence how the fibres are oriented in the paper. The induced forces will stretch and align many of the fibres along the Machine Direction (MD — the direction of material flow in a paper machine). The forces due to pressure and heat have only minor influence on the fibre orientations. The wire allows some support for the fibres while they form the network and is still sparse enough to let most of the water pass through during the initial dewatering. The press section then shapes the paper material into a more uniformly thick material and during the gradual dewatering sequence between the many heated and pressurized cylinders, the rest of the excessive water is removed. The normal moisture content of paper is a few percent. The middle part of the paper machine, the drying section (see Figure 1), is necessarily huge, since the process needs to remove water gradually to prevent ruptures in the continuous paper sheet. A *calendering* section will then im-

prove the surface smoothness. Finally, the paper is wound onto large reels that can hold several kilometres of paper, with automatic switching between reels while the machine is operating at full speed. The entire process is completely automated to enable highest possible yield, since an investment as large as this (e.g. 400 million USD) will need to run at top speed for many years to become profitable. Since the large papermaking machines are occupied by producing paper for customers, and it is critical that no tests interrupt this process, these paper machine are not suitable for research experiments. Instead smaller ones that have many options for process tuning are used. See for example the STFI EuroFEX (Mähler, 2000).

Our main tool for analyzing paper will be computerized image analysis. That is, we will develop software to enable a computer to perform the measurements. Both 2D and 3D image analysis have been used. Since the papermaking process does apply large forces on the fibres, they have a main orientation. This is roughly aligned with the MD. If we create images that are perpendicular to MD, many of the fibre cross-sections can be found. We focus on the individual fibres, but there can be other substances in paper such as: fibre fractions, filler, surface coatings, and various residues from the papermaking process. Our sample consists of fibres and fibre fractions. The following sections will describe the details of how we gradually transform a small piece of paper into a digital representation together with some examples of measurements. After the conclusions, there is a section with the published papers. As an alternative to reading the full papers, a quick glance at Section 4 will highlight the main topic of each paper.

1.2 The “3D tracking of fibres in paper” project

The goal of this project was to provide a new kind of virtual microscope, enabling users to look inside paper samples together with a computer supported ability to identify fibres and make measurements directly in the 3D structure. Especially the possibility to analyze fibres in the paper, i.e. “in situ”, would provide new measurements not previously available by methods that require the paper to be dissolved. Although our imaging technique also destroys the sample during imaging, preventing more than one imaging technique to be used, we do retain the inner structure of paper well. Alternatives, as the successful STFI Fibre Master (Karlsson and Fransson, 2000) uses 2D image analysis to gather statistical information on fibres in dissolved paper samples, and can analyze geometrical properties of thousands of fibres in a few minutes — a very efficient approach. We cannot match this fibre count, but instead provide a detailed 3D analysis of some of the fibres in a paper sample.

The project was initiated jointly by Professors Örjan Sävborg, StoraEnso Research, Gunilla Borgefors, Centre for Image Analysis (CBA), and Björn Kruse, Linköping University — Campus Norrköping, to begin a study that would investigate how paper and fibre analysis could benefit from 3D image analysis. It was called “3D tracking of fibres in paper”, indicating the need to find individual fibres in 3D images of paper, for measurement purposes. Two PhD students were assigned to the project, Mattias Aronsson (CBA) and Arash Fayyazi (Linköping University — Campus Norrköping). The rather large project was eventually divided

into two parts: (1) individual fibre segmentation, by Aronsson and (2) paper structure segmentation, by Fayyazi. The idea was that these two parts would provide complementary information about the paper structure and furthermore be combined to cover a larger aspect of paper and fibre properties. Unfortunately, due to both the complexity of this project and time constraints, we were able to fulfil only parts of this goal. As can be seen by the two theses generated from this project, we have moved the knowledge on how to analyze paper in three spatial dimensions forward and have also been able to emphasize the special demands and problems for this task. For detailed information of part-2, see the Thesis of Fayyazi (2002).

The national programme VISIT (VISual Information Technology), was initiated by Ewert Bengtsson (CBA), and one of the goals was to strengthen the cooperation between image analysis research groups in Sweden. Without the financial support from the Swedish Foundation for Strategic Research (SSF) the many interesting projects would not have been possible. See Bergholm (2002) for details of the other projects in the VISIT programme.

1.3 Paper research

Paper, although common in everyday life, is a complex material which researchers in various fields continuously try to learn more about. The practical use of paper is immense, as well as the different kinds of paper grades available. Each paper type is optimized for its intended use, e.g., tissue paper, print paper, milk cartons, and cardboard boxes. Tissue should be able to absorb a lot of fluid. Print paper should be opaque enough so that the text on the opposite side do not interfere, and also have a surface suitable for the printing technique used, enabling sharp text and images. A milk carton should be light and stiff so it does not flex too much when we hold it, have an acceptable surface to allow printing of appealing graphics so people actually buy it, and also enable an efficient packaging technique to be used. A multi-layer structure like the one used for the Skoghall board is intended to combine high stiffness with low weight. In order to be able to produce sharp 90° corners, it is also necessary that the structure can produce well defined deformations so that the folding does not produce breaks (cracks) in either of the two surfaces of the board. So manufacturing a milk-carton box is Hi-Tech! General boxes for packaging should be durable, to be able to withstand heavy loads without ruptures. There are many different research projects that deal with analysis of paper and fibres, since the process from tree to final paper is quite complex. As paper is made of wood, image analysis may be used early in the paper manufacturing process by, e.g., microwave tomography of logs (Kaestner, 2002) to automatically sort and split the logs into pieces of various qualities. By using image analysis on cross-sections, it is possible to estimate wood properties by measuring the fibre morphology, i.e., detailed shape analysis, as was done in Moëll (2001). See Figure 2 for an example of how fibres are organized within the tree. The entire path from forest to final paper spans a wide range of research topics. Paper surface treatment is important for the printing properties of paper, while fibre bonds influences the mechanical strength of paper. To understand all these aspects we need knowledge of mechanical properties, chemical surface properties, and more, to enable free fibres from the dissolved chips of wood to bond into a network with desirable properties. All these steps must be possible to

implement as an industrial process with high demands on efficiency and tolerance on the paper properties.

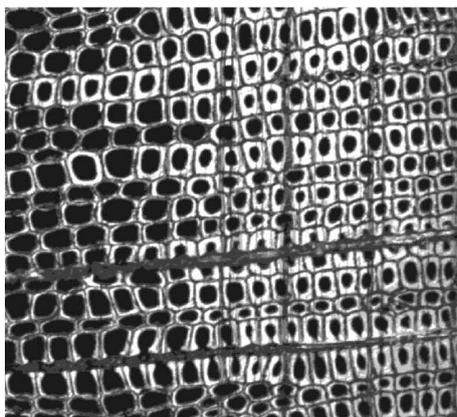


Figure 2: Transverse section from Scots pine, imaged by confocal microscopy. Size of image is $625 \times 625 \mu\text{m}^2$. From Moëll (2001).

When we started the “3D tracking of fibres in paper” project, there was only one previous attempt (Yang et al., 1978) to analyse the 3D structure of real paper samples, as far as we are aware. That project involved tracing fibre contours by hand in cross-sectional images, i.e., a small amount of computer support. Hence, the manual work needed was enormous for even tiny volumes. More recent research have digitized the 3D wood structure, i.e. to view and measure properties while the fibres are still aggregated in bundles, similar to the conditions inside trees. Using a combination of microtomy and light microscopy, Bardage (2001) and Bardage and Daniel (2002) analyzed samples of Norway spruce, regarding fibre dimensions in three spatial dimensions and differences in collapse behaviour. There have been several attempts to create and measure properties of simulated 3D paper volumes that are designed to resemble actual paper to various degrees. See, e.g. Wang and Shaler (1998) and Niskanen et al. (1997). During five years of this project, we have seen only one other project involved in recreating the 3D structure of real paper samples, performed at the Pulp and Paper Research Institute (PFI), Norway. See Section 2.4 for more details. We believe that our work can help others to understand the special problems that are important for dealing with 3D fibre analysis, and our attempts may provide a base for others interested in this topic.

Fibres provide trees with both the support and the transport of nutrients. When cells are formed inside a tree they start as cylinder-shaped liquid filled protoplasts. Then slowly cell wall material is laid down towards the outer regions and the cell is transformed into a multi-layered fibre wall, with a hole in the middle. When the cell has passed the lignification step (strengthens the cell), it resembles a tube, i.e., a hollow structure that can transport nutrition. During the lifespan of the tree, the fibre remains in this mode. Even if there are different cell types in a tree, we are only interested in those that provide the strength of paper. These cells are called fibres in hardwood (e.g. birch and willow) and tracheids in softwood (e.g. spruce and

cedar). For softwood tree species 90–95% of the tree volume consists of tracheids (fibres), the rest is parenchyma and epithelial cells. Hardwood tree species may have more types of cells. A tree is truly ingeniously “designed”, since it has the ability to grow and produce oxygen for us during its lifetime and afterwards come to use as boards for construction or paper sheets. The tree is really a very optimized production facility, and if you reflect upon this brilliant system that enables growth and transport of nutrients, it will permanently alter your mind regarding the green “stuff” outside.

Even if different tree species have large variations in fibre shape and size, a spruce or pine fibre (common in Swedish paper mills) have an approximate average length of 4 mm and have $10\text{--}60\ \mu\text{m}$ diameter, see Lundqvist (1999). Cotton and flax fibres are much longer, up to 70 mm in length. This is important, as we need to decide what sampling resolution to use when digitizing the fibres. In Figure 3 there is a set of images of the same fibre cross-section at different resolutions. By looking at some images similar to this one, we concluded that a resolution of about $1\ \mu\text{m}$ is probably necessary to have a good chance of identifying fibres.



Figure 3: Fibre cluster sampled at $0.7\ \mu\text{m}$, $1.4\ \mu\text{m}$, $2.8\ \mu\text{m}$, and $5.6\ \mu\text{m}$.

One may tend to think that $1\ \mu\text{m}$ is a very high resolution to perform analysis. It is worth mentioning that there is ongoing work to create detailed models of single fibres, i.e., submicron resolution will be needed. At the Wood Ultrastructure Research Centre (WURC) at the Swedish University of Agricultural Sciences (SLU), Jonas Brändström has worked with models for the tracheid cell wall on Norwegian spruce (Brändström, 2002). At the other end of the scale, for entire paper sheets it is necessary to measure, e.g., fluctuation in fibre distributions. An effect called *floc-culation* causes the fibres to be unevenly distributed. The measure of this unevenness is called formation. Ideally, paper should have a smooth surface with similar properties all over the paper sheet. Clusters of nearby fibres tend to form small aggregates, and these do not form into a homogeneous material as well as free fibres. To be able to measure this effect, a piece of paper with at least a few decimetres width is required.

A 3D X-ray based digitization of felt (a fibrous material like paper but often thicker and less porous) has been made by Thibault et al. (2002), but as far as we know, no further analysis of this digital sample has been made yet. Compared to our type of paper, felt is less porous, and it may be even harder to identify individual fibres in such material. The more the fibres are linked together, the harder the subdivisions will be.

Our project is between the scale of modelling a single fibre and a paper sheet. Furthermore, it is unique in the sense that we do not only suggest a viable procedure

to create a digitization of paper samples, but also have worked on methods to segment the paper into individual fibres to enable measurements on both the fibres and the network. Building a digital model of paper, it does not suffice to only digitize the paper sample. Instead, the segmentation of individual fibres will enable many more possibilities to use the digital data as a real model of the paper sample. Although we are not fully satisfied with the amount of fibres we can find, it is an important step forward.

There are other paper research projects that are more futuristic and inter-disciplinary, e.g., to develop techniques to print an electronic display on paper. This would make it possible to create a book of 100 programmable pages, add a computer, and Windows 2012 to be able to download your entire library of e-books. Maybe even read it if the Operating System boots without any errors ...

1.4 Why 3D image analysis?

Why is image analysis essential, and furthermore in three spatial dimensions? Because, for now, there are no other ways to measure properties of entire fibres in paper samples! Image analysis combines computers high processing power with the possibilities of available imaging equipments to create 3D images. The image analysis approach has at least one unique property, and that is to convert the entire geometric structure of the paper sample into a digital format, which is required in order to use a computer for the analysis. Naturally, it is not possible to capture every aspect of a complicated material such as paper, but converting the geometrical structure of paper with micrometer resolution gives a unique opportunity to do some serious “number crunching” of the paper sample properties. By realizing that the information we obtain from the digitization process is at a low-level, since it only tells us which samples on a 3D rectangular grid that are voids or not, it is easy to imagine that it requires considerable effort to find what we denote as “fibres”, which is a high-level description of the objects we are interested in.

We know that fibres are structures that are distributed inside a 3D structure (the paper), so using any lower dimension than three, will result in loss of vital information. For *some* properties, e.g., fibre diameter, you may assume that the fibre has the same diameter along its length, and use 2D image analysis instead. However, this is a much less accurate approach, than actually measuring the diameter along the whole length to verify whether the width changes. The attainable precision of our approach is thus very high, and if we are able to minimize the artefacts introduced by digitization and software, an estimate of the absolute accuracy would be possible. For now, we have focused on the implementation of the reconstruction and fibre segmentation, without exactly verifying the absolute accuracy attainable. This could be troublesome, since it may require an alternative technique to be absolutely sure about the accuracy. Nonetheless, 3D image analysis of paper definitely provides a new higher level of detail, when studying the internal structure of paper. See Figure 4.

Generally, 3D analysis tends to be much more complex than 2D analysis of images. Thus, it is seldom a good idea to use 3D images when 2D images would

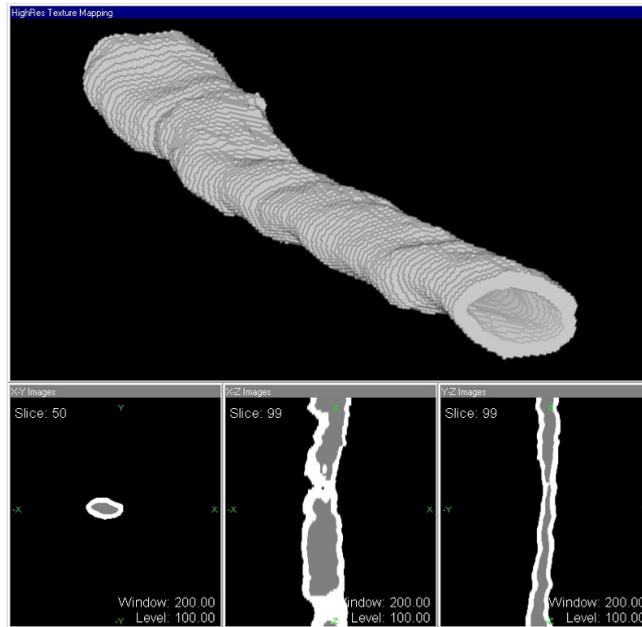


Figure 4: Rendering of segmented fibre, using the technique in Section 3.8.3.

suffice. The basic difference between 2D and 3D images is that a unique position in each image is specified by (x, y) and (x, y, z) , thus two or three integers are needed. See Figure 5 for an example of images stored as rectangular/cubic (2D/3D) grids, the most common format. There do exist other formats, e.g., hexagonal grids, which has some advantages, e.g., that a single connectivity (and distance) of six local neighbours for both the foreground and background regions can be used. In square and cubic grids it is important to use different connectivity for foreground and background, see Rosenfeld (1979) for details. Traditionally, each sampled value in a 2D image is called *pixel* and in a 3D image a *voxel*.

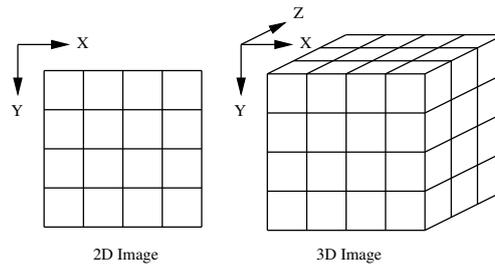


Figure 5: Principal difference between 2D and 3D images is the number of coordinates needed to access a specific pixel/voxel.

As you will see, even if we have had limited success in recreating the entire fibre network, none of the 3D measures on the individual fibres would be possible if

only 2D images were available. When an improved segmentation method becomes available, the fibre measures we have developed can benefit from this, and some of the properties of the network structure (bond area, mean free path, orientation distributions, etc.) will be possible to estimate. Thus, the key to unlocking the internal paper network structure is an accurate segmentation of many of the individual fibres, a much more intricate task than the simple segmentation that differentiates fibre material and void.

The processing power of computers nowadays allows more advanced computations, so three spatial dimensions are used in a variety of applications. One of the areas where 3D image analysis really has displayed some of its potential is in medical imaging (Udupa and Herman, 1999), where it is now commonplace to create 3D images of the brain and internal structures of the human body. Without proper software, some of the imaging equipment used as, e.g., CT, PET, SPECT and MR, which all are ingeniously developed, would not be as useful as they are now.

1.5 Digital 3D model of paper

Simplification of a problem to a mathematical model that captures the most important properties is often required. Even with today's computers that are able to perform blazingly fast calculations, the computational budget will always be a limiting factor, so it is important to not overdo the amount of detail in the model. By using a 3D voxel model, we sample the paper at a 3D rectangular grid and provide at each sample point (voxel) a value of the estimated mean atomic weight. This can be used to differentiate between fibre materials and void (pockets), see Figure 7. Because of the large amount of voxels needed to represent even a tiny paper sample, the main task is to transform this low-level information into a high-level one, i.e. individual fibres. See example of a voxel model of a single fibre in Figure 6.

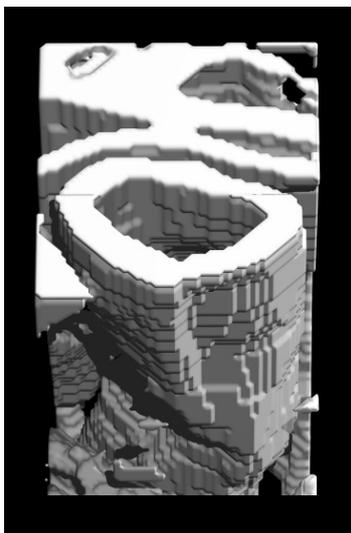


Figure 6: Example of a part of a single fibre, as a voxel model.

We thus model the object in all three spatial dimensions but allow only a fixed level of detail, using a smallest unit, the voxel. This creates volumes with a bit of “Lego-ish”-style when rendered at high magnification, but simplifies the processing, since many subtasks only require local calculations to manipulate the volume, e.g. filtering etc. One important topic is how to choose the size of the voxels. The best and simplest way is to use cubic voxels as then the resolution is equal along all three dimensions. This is often not possible, due to the acquisition process. Then we have to take the different resolutions into account during the calculations. Less resolution along one co-ordinate axis will mean that we have less information about the geometry in that particular direction. It is not possible to fill in that missing information, but we can at least assume locally smooth behaviour of the fibres and interpolate, which is a small improvement that will attenuate the artificial jaggedness introduced by the digitization process. Due to the relatively sparse sampling we cannot use simple interpolation, but need more advanced methods.

We use $0.7 \times 0.7 \times 5 \mu\text{m}^3$ parallelepiped shaped voxels, as a compromise between getting high enough resolution to be able to find the fibres and as few voxels possible to be able to minimize the memory requirements. Due to the imaging process the voxels are not cubic and this is one of many complications for us.

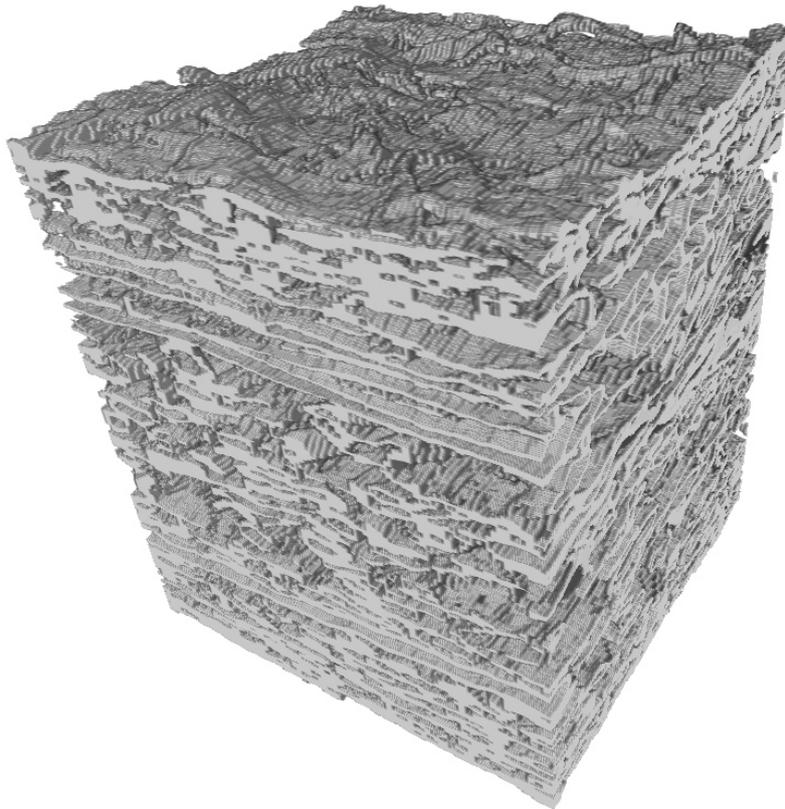


Figure 7: Voxel model of a milk-carton paper sample, size $500 \times 768 \times 102$ voxels.

2 Data acquisition

Basically there exist two main approaches to create digital volumes: (1) tomography and (2) serial sectioning by microtome cutting, to create stacks of 2D images. We were fortunate enough to receive several data sets, but due to quality and time limitations we have concentrated on one specific set, that we believe had the greatest potential for image analysis. Örjan Sävborg and Olle Henningsson at the StoraEnso Research facility in Falun, put much effort into making this data set. We are really happy about these data, as such a unique data set provides an excellent opportunity to test the robustness and accuracy of developed methods on a real paper sample.

2.1 Sample preparation and scanning

Many of the best imaging devices available cannot capture 3D images, so we need to obtain a series of 2D images instead. The machine that makes this possible is a microtome, i.e., a diamond (or glass) blade equipped high precision cutter. See Figure 8. This high-tech “cheese cutter” give extremely thin, in our case $1\ \mu\text{m}$ thin slices. Special “Ultra Microtomes” can cut even thinner sections down to $50\ \text{nm}$. There are other techniques, such as laser evaporation that “boils off” thin layers of material, but they may produce surfaces that are uneven and this is not good enough for our purpose. Some examples of imaging devices are given below.

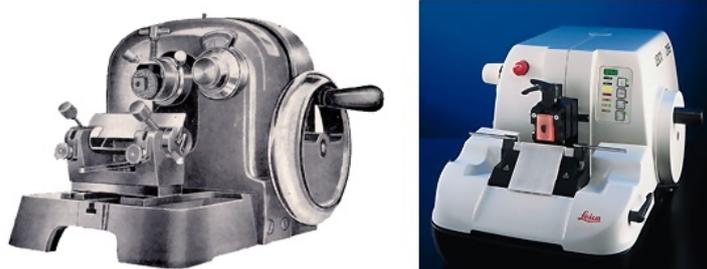


Figure 8: Example of two microtomes. Even if there are quite some years between these two models, they look rather similar.

2.2 Light Microscopy

The first data set we had available was kindly provided by the Swedish Pulp and Paper Research Institute (STFI, 2002) and prepared by Joanna Hornatowska. A small digitized paper sample was created, by a combination of slicing off thin $1\ \mu\text{m}$ slices of paper embedded in plastic and manually placing these on object glasses, then staining with toluene blue to enhance the contrast before imaging with a CCD (Charge Coupled Device) camera. See Figure 9 for an example. A good property of this set is its high dynamics, since the optical system can capture a wide range of opacity through the thin samples. It is a delicate manual process to slice off micrometer thin sections and manually put these on object glasses. The amount of

geometrical distortion, or even loss of some slices, introduced by this process, requires a complex registration procedure, i.e. geometrical correction. We tried an approach based on vector displacement fields, where you manually identify corresponding “landmarks” in two adjacent images. This will create a set of 2D displacement vectors at the landmark co-ordinates that can be interpolated into a continuous 2D displacement field to correct the entire image. Even if the distortions were attenuated, it was not good enough to be useful for a complete volume reconstruction. At that time, when we were developing and testing this approach, we did not realize how complicated this would be, a kind of “hubris” that led to one or two statements of the kind “this will soon be solved”. It is easy to believe that just because you can easily identify a fibre using your eyes, it would only be a matter of coding to let the computer do the same. However, some hands-on experience will tell you that few image analysis problems are trivial.

Confocal microscopes are a special type of light microscopes that has an optical slicing capability. The lens system uses a pinhole to create a very thin focal plane, this makes it possible to collect light at a well defined depth. More porous paper grades than milk cartons, that we are interested in, may be well suited for confocal light microscopes.

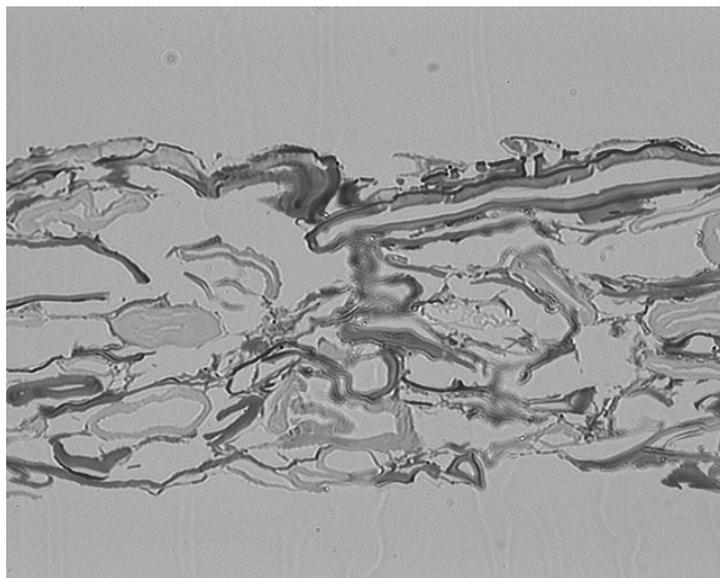


Figure 9: Light microscopy image of LWC paper cross-section. Prepared and imaged by Joanna Hornatowska at STFI.

2.3 Scanning Electron Microscopy

Scanning electron microscopy (SEM) has been used for many imaging applications. Most people have seen the images of a CPU chip or head of a fly in books, with an amazing resolution of small details, far better than any optical system based on

visible light can approach. See Figure 10 and Figure 11 for examples of a circuit board and a tiny knot of a hair at high magnification.



Figure 10: LEO SEM image of small part of a printed circuit board, including surface bonded IC and resistors. Copyright © LEO Electron Microscopy Ltd.



Figure 11: LEO SEM image of a human hair, tied into a tiny knot. Just as in the Wella shampoo commercials you can see the surface breaking up a bit. Copyright © LEO Electron Microscopy Ltd.

A modern SEM, as the LEO 435VP that was used here, has an ultimate resolution in the nanometre scale for an ideal test case. Our use of SEM to cover fibre details around $1\ \mu\text{m}$, thus by no way stresses the equipment to the limit. Using back-scatter mode, which approximately estimates sample atomic density, we were able

to use the paper sample without any preparations except for embedding in epoxy, to make a stable block of material. The rather compact desktop SEM device, see Figure 12, consist of a camera module, a computer (Windows based PC), two monitors and an input devices.

Since SEM is not a true 3D imaging device, we need the microtome to slice off and capture the inner parts of the paper. One crucial insight to fibre imaging was that it is important to use the ability of the SEM camera to image the surface of the remaining paper sample, instead of the thin section sliced off during each iteration of imaging. This approach will result in much lower geometrical distortions, as the paper surface we image has a good support. This is a difference compared to LM, where each slice was imaged by collecting the light that pass through it. The level of contrast possible is mainly a question of acquisition time, but as we need to make many 2D images and also slice the sample between each imaging session, it was important to minimizing the time for each slice. Selecting backscatter mode, results in nearly binary images, which is an advantage for us, since we are only interested in differentiating between fibre-wall material and void in this application. To improve the discrimination ability even further, we could have increased the acquisition time to raise the contrast and simultaneously attenuate the background noise.

2.4 X-ray micro tomography

At Trondheim University in Norway, Emil Samuelsen (NTNU), Rune Holmstad, Christine Antoine, Per Nygård and Torbjørn Helle, at the Norwegian Pulp and Paper Research Institute (PFI), have created a 3D model of paper by using the European Synchrotron Radiation Facility (ESRF, 2002) in Grenoble, France. To get sufficient resolution, they used phase-contrast mode, which locates boundaries between fibre material and void. Using image analysis they could convert back to the preferred filled fibre “mode”. Some issues of artefacts due to the tomography approach (Herman, 1980), are problems that needed attention. This data set was made available to us just recently and we have unfortunately not been able to explore it as much as we would have wanted. It is interesting to compare various existing techniques to acquire 3D images from paper samples, since the imaging process has a large impact of how well it is possible to analyze the data. For more details on this approach, see Samuelsen et al. (1999).

2.5 Combining micro tomography and SEM

We decided to combine a Scanning Electron Microscope (SEM) with a microtome, to create a stack of 2D images that we later will assemble into a 3D image. Cutting paper with this very sharp knife does expose the knife to high pressures, since the area of contact is so small. That is why we make five $1\ \mu\text{m}$ cuts instead of one $5\ \mu\text{m}$ cut between each imaging step, for the SEM setup we describe below. Another trick that can be used is to add water, which let the thin paper section we cut off float on top of the water which can prevent the slice to break apart during cutting. The SEM camera is equipped with a computer for the imaging and parameter setup, and was also used to store the captured images. The process, which requires some delicate manual manipulation, is described below:



Figure 12: The LEO Scanning Electron Microscope we used. Copyright © LEO Electron Microscopy Ltd.

- Embed the sample in Epoxy resin, together with Rayon reference wires at the upper and lower surfaces
- Repeat N times
 - Capture cut-surface with M overlapping 2D images, by manually repositioning the sample in the SEM camera
 - Repeat $M - 1$ times
 - Move the SEM camera capture area 80–90% of the image width
 - Capture cut-surface
 - End
 - Use microtome to slice off $5 \mu m$ of the outermost layer, by a sequence of five $1 \mu m$ cuts.
 - End
- End

After this process, a set of 2D images that captures the internal structure is the result, see Figure 13.

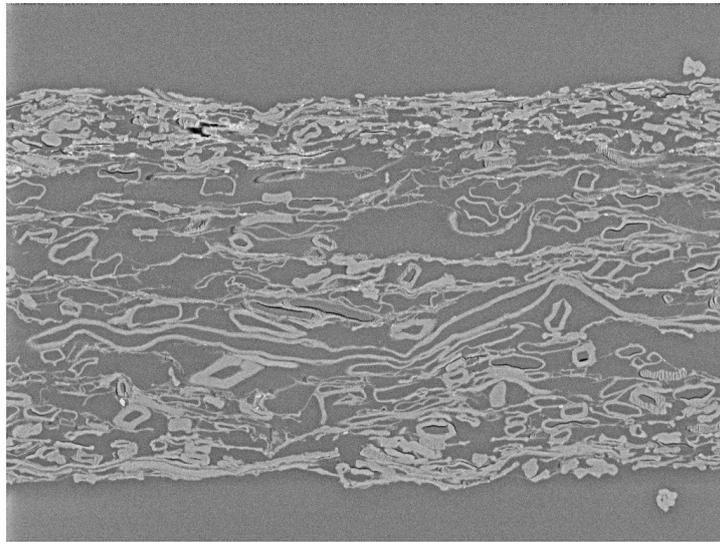


Figure 13: SEM image of milk-carton paper cross-section. Prepared and imaged by Örjan Sävborg and Olle Henningsson at StoraEnso Research.

2.6 Imaging conclusions

The SEM/microtome setup seems to be the best choice, when we compare the pros. and cons. of the three imaging setups in Sections 2.2, 2.3, and 2.4. Most important benefits are the level of detail, ease of first level segmentation between fibre and void, and low geometrical distortions. However, X-ray micro-tomography has some important benefits, mainly because two issues are solved by the X-ray approach: (1) the tedious manual sectioning of the volume and (2) the need of registration is eliminated. Tomography manufacturers like Siemens, Philips and LEO, keep improving the resolution and contrast, so this technique has great future potential to be a good alternative to slicing based approaches. The level of automation possible is much higher, since the slicing is performed “virtually” by sending a thin fan of X-rays through the sample and rotate the sample between each image capture. However, the artefacts caused by the tomography process also need to be attenuated to enable the delicate fibre segmentation. Optical microscopes can be used with great success, performing 2D image analysis. But they seem less suitable for volume reconstructions, since they generate images with large distortions and also the segmentation, even only to separate between fibre and void, is much more difficult to carry out in these images. Note that when we say large distortions, we mean something above 0.01 mm . In other applications this may be acceptable, while it is far too large for the average fibre diameter of $0.01\text{--}0.05\text{ mm}$. There are really many possibilities to combine different devices to acquire the images of interest. Confocal microscopes have been used to create 3D volumes, although as many paper grades are dense and quite thick, it will not be possible to image very deep through the surface of dense paper grades. Instead, a combination of sectioning by a microtome and imaging with a confocal microscopy could be an alternative technique to create digital 3D

paper samples. This decreases the amount of sectioning needed, as the confocal microscope can create several images at different depths between each sectioning pass.

There are many different 3D projects currently in progress, which use different types of imaging techniques to digitize as diverse things as viruses, cells, beetles, frogs, humans and more. In addition to the new knowledge this can eventually provide, it is also great fun just watching the images. See Figure 14 for an example of a stag beetle, imaged using X-ray micro tomography.

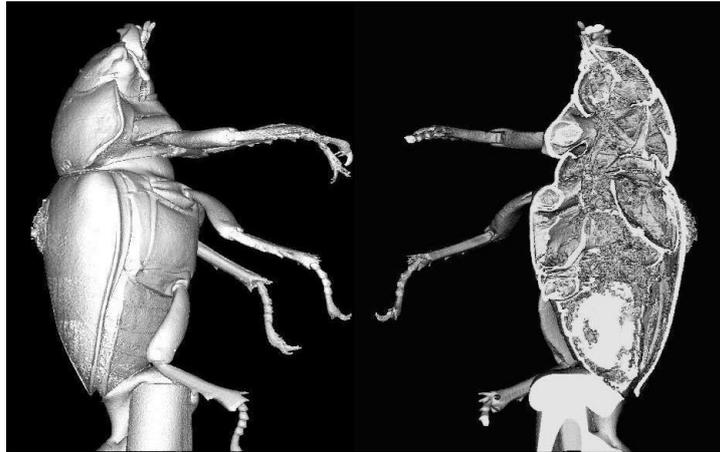


Figure 14: An X-rayed beetle, voxelsize $80\ \mu\text{m}$. Copyright © Dr Graham R Davis at Queen Mary, University of London, UK.

3 Methods

In this section, we first give a general introduction about some of the specific issues about 3D image analysis that we believe is important to keep in mind when you read the papers included in this Thesis. Thereafter we describe the process we have developed to transform paper samples into digital models, together with some measurements.

3.1 Implementing efficient 3D image analysis algorithms

It is a huge leap in storage requirements, performing analysis in three spatial dimensions instead of two spatial dimensions. Efficient calculations of large data sets will require methods that do not need the entire 3D image to be loaded into the memory of the computer (RAM — the Random Access Memory, used for fast access to data). Redesigning methods so that efficient load/store of subparts of the image is possible during calculations presents quite a challenge by itself. We have not focused our methods on this aspect. Analysis of 3D images has much to benefit from parallel processing techniques, since many problems need massive sets of computations. The additional problem is then to develop image analysis algorithms suitable for parallel execution, that speed up calculations by solving independent subtasks on different Central Processing Units (CPU, e.g., Intel Pentium-4 and AMD Athlon) and combining the results. As 3D images requires much more storage space, it is a good idea to add various compression techniques that are suitable for fibre images. Using a high-resolution grey-level 2D image means something like a 1024×1024 pixel image, requiring 1 Mbyte of storage space in RAM. If we need the same resolution for measures along three co-ordinate axes, we have a 3D image consisting of $1024 \times 1024 \times 1024$ voxels and need 1 Gbyte of storage. Disc space for file storage is seldom a limiting factor nowadays, since 80 Gbyte drives are considered the default purchase. Even if RAM for computers have shown a similar impressive increase in capacity, modern computers have 128–1024 Mbyte of RAM. This is used for the various programs that the operating system consists of, a file cache to speed up hard disc access and the rest for the applications, e.g., our fibre segmentation software. Simple things, as storing intermediate results, can be a real bottle-neck working with 3D images. Therefore, it is important to design with data size in mind, as well as efficiency of calculations. For 3D images that cannot be held entirely in RAM, the disc access pattern from the algorithms also need to be optimized with respect to the file cache system in the operating system, otherwise severe performance degradation will occur. In addition to the complexity of algorithmic design, also some computer science aspects have to be considered for optimal performance. Each of these issues is rather complex, so we have not incorporated them as much as we would have liked. More optimization according to above issues would definitely increase the possible size of the paper sample, as well as decrease the time of computation.

Many of our methods are based on 2D/3D distance transformations on an image, which will be denoted as DT , and $DT(\text{parameters})$ when the specific type is of importance. A DT has one image as parameter, and the result is also an image,

which we denote as DT_{image} . For details about DTs, see Borgefors (1986, 1994, 1996). A DT can be used for N:th dimensional images, but we use 2D and 3D DTs. Below we assume 3D images, thus *voxels* are used. The DT measures for each voxel its shortest distance to the voxel set initiated to zero before the calculation (all other voxels are initialized to the highest possible value, as we mostly use an 8-bit unsigned byte data type, this is 255, but in principle ∞). There are different types of DTs, but we will always use the Chamfer based DT with integer parameters optimized for minimum error and maximum coverage using the small unsigned byte data type, as we do not need exact distance measures, instead favor fastest possible calculations. The Chamfer DT can be calculated very fast by small local masks, that scan two times through the image/volume. If very exact measures are needed, a chamfer euclidean transform can be used (Danielsson, 1980). This method takes more time but have a maximum error of only 0.09 pixels, still using a fixed number of scans (four needed).

As measuring distances is a very common task, we need an efficient method for that, and the Chamfer DT provides that. Sometimes we need to also keep the information from which region we measure distance, then a *labelled DT* is used. Also *reverse DTs* are used. These are similar to implement and are useful for spreading a discrete set of non-zero valued pixels/voxels, the distance of spreading based on the initial values. Additionally, some cases require that the DT is measured as the shortest path along restricted set of voxels. This DT is called a *constrained DT*. A nice example is a maze, where you need to find the shortest path from start to exit, and are not allowed to move through walls. Instead of using a graph search method as depth first search, a restricted DT will calculate the path for you. One disadvantage of restricted DTs is that the calculation time is dependent on the image contents, which is not the case for basic DTs. Thus, an sequence of chamfer updates are needed until no change in the DT_{image} occurs. An example of a worst-case scenario is a spiral shaped figure. Apparently, there are many different types of distance measures in digital images, and we will make heavy use of these methods.

A simple compression technique, that we experimented with, is the Run Length Encoding (RLE). RLE is computationally efficient, simple to implement and can provide an algorithm with streaming data, i.e., you can decompress “on-the-fly” and do not need to decompress the entire volume into RAM at once, see Figure 15. For the common tasks of looping through the 3D image in a fixed $x-y-z$ order, we can compress in that same order. For the binarized volume we obtained compression ratios around 20, using RLE. For example, instead of needing 230 Mbyte, 12 Mbyte is enough. Other techniques, e.g., Octrees (Wilhelms and Gelder, 1990), may be more suitable for algorithms that require 3D sub-blocks of the image in each step.

Another type of compression is how we encode values. For example, during rendering, we need to store the normal of each voxel to be able to create realistic images of the model. It is not a good idea to store this 3D vector as a triple float data type, since this requires 24 bytes of memory for each voxel. Instead, we need to encode the discrete set of directions with as few bits as possible, and can take advantage of the fact that normal vectors are unit-length. It is sufficient to store the direction using spherical co-ordinates. Allowing, e.g., 256 different directions along each co-ordinate in the (ϕ, θ) system, it is enough using only two bytes for

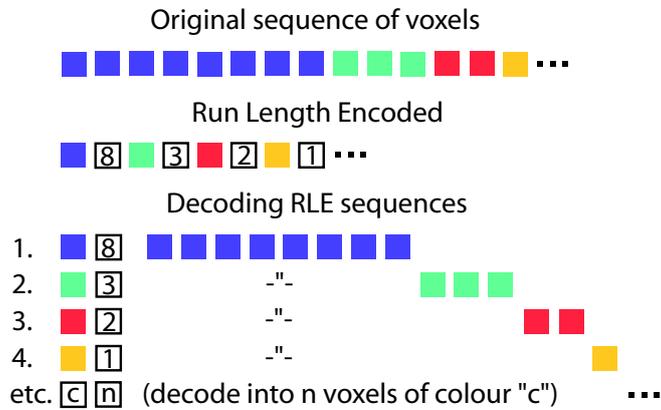


Figure 15: Run Length Encoding is a general compression technique. Even if it mostly compresses less than, e.g., bzip2, its strength is the efficient encode/decode process.

each normal vector, resulting in a compression about 12 times.

A voxel is our smallest building block, similar to a physical atom model of an arbitrary object. There is a huge number of voxels in many 3D images, and this makes it necessary to be extremely restrictive on the number of bits for each voxel property we need to store. Compare to, e.g., CAD-like surface structure models with spline-based patches. As each of these can cover a larger piece of the object, it is possible to use more bytes for each patch, but even these need to be optimized for size. For example, an accurate CAD model of a car engine requires millions of triangular patches for each part with complex non-smooth shape. The preferable digital representation of an object is dependent on what object we want to digitize and to what degree of geometrical accuracy. For fibres, we are confident that the voxel based model is optimal, since the representation is very compact, considering the amount of sample points from the SEM camera we have. One of the properties that makes voxel images compact, is the 3D rectangular grid restriction, which enable the use of implicit (x, y, z) co-ordinates for the centre position of each voxel.

We are using a combination of MATLAB[™], C (stand alone as well as for the MATLAB MEX-files and the IMP program, see Nordin (1997)) and C++ code to achieve reasonable computation times for the methods used. For massive computations it is important to choose a language that compiles the code and since C/C++ is well known for its efficiency, we use these. Since C++ do not have basic matrix calculus, we use the Template Numerical Toolkit (TNT) package (Pozo, 2002) to be able to write code with vector and matrix calculations. The Standard Template Library (STL) for some of the data types is needed. A basic 3D image class data type has been implemented to allow a common base for coding, load/store image, simple algorithms, etc. You may download the software from

3.2 Raw data

We want our captured 2D images to each lie in a (x, y) co-ordinate system, with the third axis, z , along the direction of microtome slicing. Contrary to common use in paper science, the paper surface will lie in the (x, z) co-ordinate system. Thus, the y -axis is approximately normal to the surface, not the z -axis. See Figure 16.

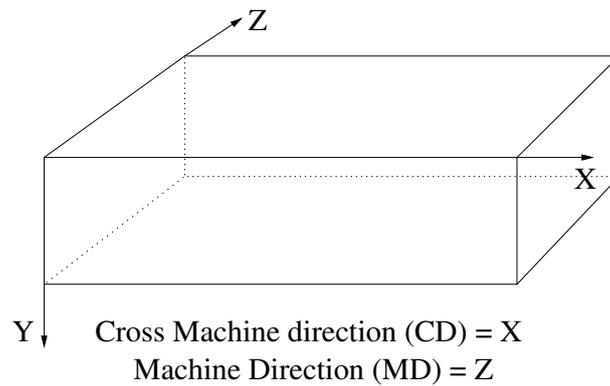


Figure 16: This is how we define the co-ordinate axes.

We have a piece of paper that we want to digitize, i.e. convert into a set of digital samples evenly spread throughout the paper. As we concluded earlier, the SEM camera will be used, which is a 2D image capturing device. We combine this camera with a microtome to be able to capture the inner parts of the paper sample. A sequence of $M \times N$ images, will be used as raw image data for the volume reconstruction. N is the number of z slices, determined by how many times the microtome is used. M is the number of overlapping images for each z slice, used to get a wider field of view than 1024×768 , which the sensor in the SEM camera provides. Since we used three overlapping images ($M = 3$) and 102 slices ($N = 102$), we have a set of 306 images. The following subsections will fill in the details about: registering each triplet of images into a set of wide images; register wide images by pair-wise comparison using landmarks to form a 3D image; shading correction and filtering to enable the use of a global threshold as binarization. Binarization is the first level of segmentation — differentiating background and fibre material.

3.3 Registering the set of images into a 3D image

Since we have to split our paper sample into a series of 2D digital images to be able to image the internal parts, we need to reassemble these images into a 3D image. We unavoidably will deform the paper sample during cutting, and these geometrical distortions have to be attenuated even though they are small. There are other distortions during imaging as well, but the geometrical distortion is the most

prominent one. Non-linearities in the response of the sensor are not a problem for us, since we will binarize the sample.

The registration process is divided into two parts. The first is to align and combine the M overlapping sub-images (i.e., images such as in Figure 13) into one wide image, see Figure 17.

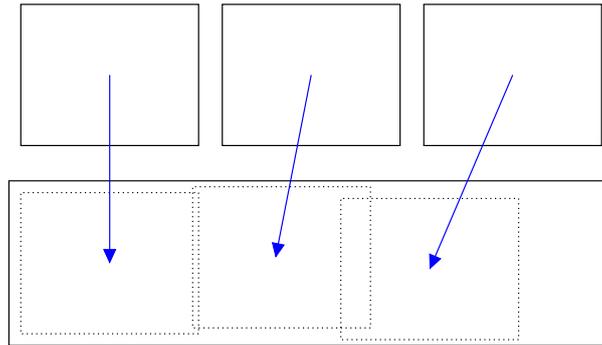


Figure 17: Each triplet of images needs to be registered before mosaiced into a wide image.

We do not know the exact amount of translation in the moving-step during digitization, so an automatic matching technique is used. We selected the Hierarchical Chamfer Matching (HCMA), see Borgefors (1988), since it is a fast and robust method for image registration. Hierarchical means that the method sub-samples the image to a hierarchy of resolution levels (Sonka et al., 1999) resembling a pyramid. The advantage of this process is that the matching can start at a lower resolution and continue further down to higher resolution levels until the original image is used for the final matching. In each step a local surrounding of $p \times p$ pixels are used for the comparison, and as the resolution in higher levels are lower, we take large steps for each local step. Ideally, this should result in both improved robustness and a much faster method, compared to standard correlation applied directly to pairs of images. In this case, correlation is measured as follows. Both images are transformed by a Canny edge detector (Canny, 1986), and one of them further transformed using a 2D DT(3,4). Translation within a fixed local surrounding and integrating the square values in the edge-pixels that are covered by the image we match against, estimate how good the match is. For best match, it is reasonable to optimize for maximum number of overlapping edge voxels and the ones that do not overlap should be as close as possible, i.e. have small DT values. To increase the robustness even further, we use standard correlation in a small local surrounding of the translation (dx, dy) that HCMA found. If HCMA was successful, we will see a local maxima in the correlation at $(0, 0)$ and lower correlation values around this local correlation maxima. This registration is not so complex, since a subset of pixels is present in the adjacent images. However, whatever your favourite matching method may be, the overlap does vary a lot, so the method you choose should be robust enough to handle this. The HCMA plus correlation works well, see Figure 18. When studying the result image after transformation and the mosaicing, it is often impossible to see the joint

between the sub-images, even without any grey-level interpolation in the transition area. This registration process is completely automatic.

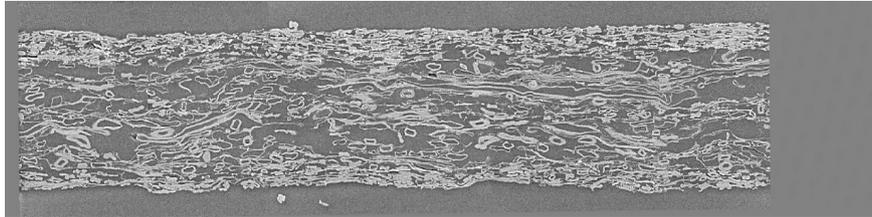


Figure 18: Example of registered and mosaiced wide image.

The second part is to assemble the set of wide images into a 3D image. Due to the $5\ \mu\text{m}$ distance between consecutive images, they differ considerably. Thus, a simple global pixel based comparison, since e.g., correlation based matching, will not work. There exist more advanced versions of correlation based matching. We tried the method in Ourselin et al. (2000), developed for rat brain cryo-sections. It uses a hierarchical approach with a set of local correlation masks to estimate a deformation field. The local masks are relatively small, so even small rotations can be captured by only translating the masks around a local neighbourhood. Then, for the point matching for the centres of the matched blocks, M-estimators are used since they are less sensitive to outliers, known to cause troubles for Least Squares (LS) based methods (Björck, 1996). The difference between M-estimators and LS methods is the minimization function. In LS, we use $\min(\sum r_i^2)$, i.e. we try to minimize the sum of the square residuals (errors). As LS assumes that all data points are equally valid, a single outlier can result in a very inaccurate solution. In M-estimators, we replace the square function by $\rho(r_i)$, i.e. $\min(\sum \rho(r_i))$ and use e.g., $\rho(r_i) = \|r_i\|_2$. By using a function ρ which does not square the residual, outliers will not influence the calculations as much. We still need a way to solve this estimation problem, and it is common to reformulate the equation above as a sequence of weighted LS problems: $\min(\sum w(r_i^{k-1})(r_i^k)^2)$. In doing so, we can essentially use the same calculations as in LS, but have a much more outlier insensitive method. The hierarchical idea is similar to HCMA, for the purpose of speed and robustness. Although the initial tests seemed promising, the method sometimes diverges and produces bad registrations. Since we had an alternative, based on landmarks, we decided to use that instead. Robustness to outliers is essential in landmark based registration. If good landmarks can be embedded into the sample, this will yield both a faster and a more robust method than image based registration. If we have the special knowledge of certain corresponding points, it is good to make use of this information.

Landmark based and image based matching are rather different. In landmark based matching, you select a set of points that you are confident can be tracked in the images to be registered. No other pixels are used for the registration. A commonly used idea is to limit the type of movements that can be corrected. By imposing this restriction, we get a simpler well defined parameter estimation problem. Our images have basically four types of movements/deformations: translation, rotation,

scaling, and a small amount of shear (the last caused by slightly uneven cutting). See Figure 19. An affine deformation model is ideal for this case. The parameters for the affine transform are estimated and if the landmarks are well chosen and the transform accurately models the deformations that occurred, the residuals should be small. The parameters, together with the transformation formula selected, are applied to all pixels, using, e.g., bilinear interpolation, since transformed pixels will generally not be aligned to the rectangular grid. It is important to use robust methods, because if a single landmark is severely misplaced, it should not be allowed to ruin the match.

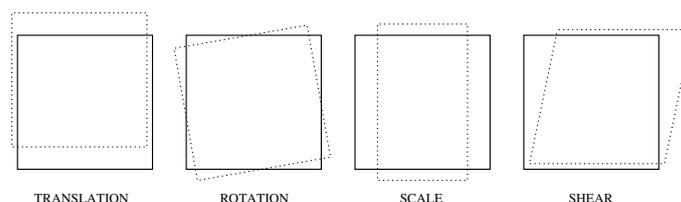


Figure 19: The four types of movements/deformations we need to undo, by registering the sub-images into the common co-ordinate system of the wide image.

The technique of image based matching does not utilize any special knowledge of points to track, but instead tries to compare sub-images directly. Basically, each sub-image is considered a landmark, but since these are scattered automatically over the image, they are easier to obtain than landmarks. The disadvantage is that the difference between a good match and a random match can be small, resulting in many false matches. However, if it works, there is no need for finding good landmarks, which can be troublesome.

As landmarks, we manually selected a bounding box for each of the two rayon threads, and a few thick-walled uncollapsed robust fibres that are found in both images. Then the centres of gravity were used as landmark points. This pair-wise matching can lead to a sheared volume, which would need to be corrected, but since our volume “only” consists of 102 slices, we did not notice any severe shear effect. It is important to ensure that the two rayon threads are straight, otherwise the errors in the pair-wise matching could accumulate to produce a volume which is not a parallelepiped. This was achieved by tiny clips at each end during curing in the epoxy resin compound. Since these rayon threads were fragile, it was quite tricky to apply enough force to hold them stretched without breaking them. We only managed to add two. It would have been desirable to have more, since then we would not have needed any additional fibres for the landmark matching.

The Least Squares (LS) method was accurate enough for our purposes. We also tried weighted LS to be able to assign different initial levels of trust to each landmark and then have an update based on each residual in an iterative manner. Since the improvement was insignificant, the standard LS method was used. It is always better to have data with less distortion, that enable simple methods, than bad data requiring advanced methods to try to recover the initial shape. There are several alternative ways to perform this registration, since slicing and reassembly is

common in medical image analysis. Since we were satisfied about the performance of our rayon thread and LS approach, there was little reason to try anything else.

3.4 Filtering and pre-processing

By analyzing the histograms from the captured 2D images, it is possible to identify two approximately Gaussian distributions that are mixed but still quite separable, see Figure 20. These correspond to signals from the resin filled areas outside the paper and inside the paper, fibre lumens, void regions, and the fibre material. The images have an uneven background, which was attenuated by using a new DT-based shading correction (see Paper II). Some of the images have a shading problem at the left border (caused by the SEM camera). It was necessary to remove this artifact too, as described in Paper II. The definite advantage of attenuating artefacts as much as possible before attempting to binarize the image, is that it is sometimes possible to use a basic segmentation method after this pre-processing. Without the correction, it would not be possible to get an accurate segmentation by a global threshold. Binarization in itself can be a very difficult step, but since this partitioning is often crucial for all of the following steps, it is wise to spend some effort on it. It is possible to not use a definite binarization partitioning into background/object, but instead to use a fuzzy approach (Udupa and Samarasekera, 1996; Borgefors and Svensson, 2002) of assigning a 0–100% degree of belongingness to the object and background for each voxel. This may aid in following individual fibre segmentation steps, but since we believe this to be part of the fine-tuning and not critical for the overall performance, we postpone it to future studies.

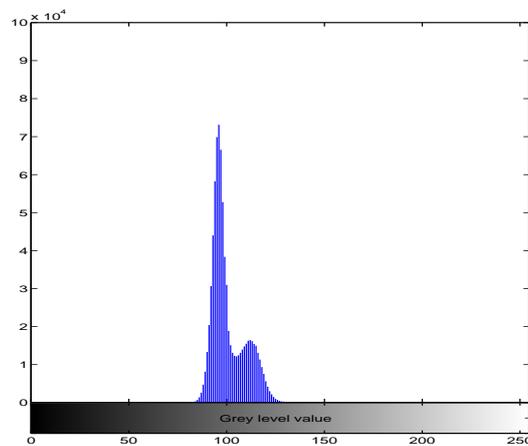


Figure 20: Example of histogram of a 2D grey-level image from the data set.

Noise in SEM images is mainly influenced by the acquisition time. Longer acquisition time mean that more electrons reach the detector in the SEM device and compared to the random electron noise count, we can get a high signal-to-noise ratio (SNR). We want to minimize the time of digitization, and we do this by reducing the

time to about 35 seconds for each capture. This produces images with a reasonable SNR. A classical image analysis problem is to remove as much noise as possible, while trying not to blur the edges in the image, which are conflicting goals. We use a median filter and an anisotropic filter based on a 2D DT, described in Paper II. This step does not seem crucial. It is possible to select your favourite method for image improvements, such as to attenuate Salt- and Pepper noise. Still, it is important to keep as much information as possible to facilitate segmentation, which is the complex part of the problem. We neither want crucial details lost due to too much smoothing, nor any scattered pixels left due to noise.

3.5 Resolution enhancements

To partly compensate for the lower resolution along one axis, we tried a linear shape-based interpolation method (Udupa and Raya, 1990) on a small piece of one fibre (see Figure 21). Simple grey-level interpolation is not enough to compensate for the relatively large displacement of fibres that exist. Instead shape-based interpolation (also known as “Morphing”), which is much more complex, is needed. Basically, we want to track the fibre path in three spatial dimensions, and predict the path where we have no samples. The method is special, since it can interpolate without exact segmentation of individual fibres, within reasonable levels of fibre displacements and deformations in shape. Still, for more complicated parts, the method has problems and needs guidance of what to interpolate. As we have not yet segmented the fibres, it is hard to provide this information. The anisotropic data may look like a set of stacked slices when rendered, but it is sometimes possible to improve the visual appearance considerably, by using shape based interpolation, see Figure 21. However, interpolating the entire 230 Mbyte dataset into isotropic resolution (same resolution along all three co-ordinate axes) would result in 1.6 Gbyte of data. Apart from the increase in storage, you have an insecurity during the interpolation step, that you actually may join parts that should have been left separate. Essentially, there is no ideal substitute for the missing information, but in some cases shape-based interpolation can be the solution to intelligent interpolation, when grey-level interpolation is not enough, which it seldom is.

3.6 Where is the surface?

Maybe surprisingly, one troublesome paper measure is to define where the upper and lower paper surfaces are located. Since the network of fibres will create a web of voids, the transition is gradual, from “definitely inside” the paper to “definitely outside”. The problem with finding the paper surface is that paper can have rather large holes and porous regions and tunnels into the inner parts, a robust paper surface detection method should not go into these pores. Instead it should capture the general trend, but at the same time not smooth too much since that would decrease the level of surface detail.

To actually calculate these two surfaces in three spatial dimensions for a 230 M byte dataset can be cumbersome, but we use a very efficient 3D rolling ball algorithm (Frank, 2000), that only requires two 3D DTs and a simple global threshold

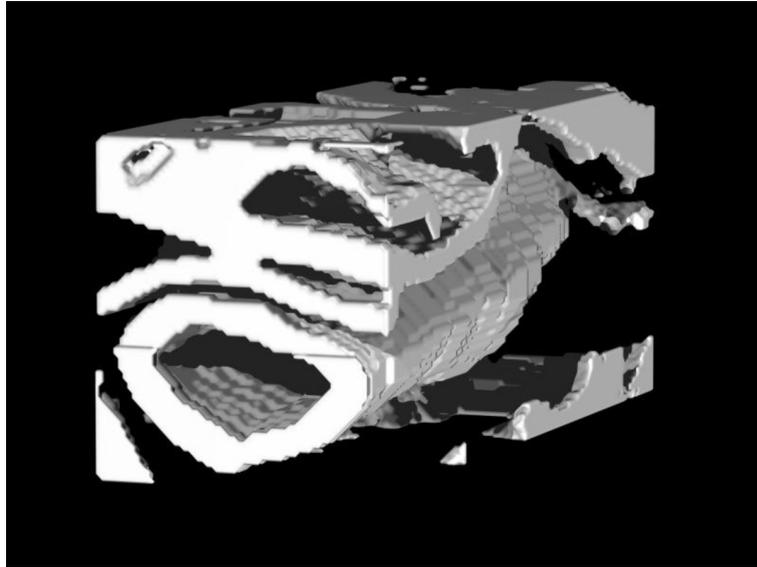


Figure 21: By using shape-based interpolation, the aliasing caused by severe under-sampling can be attenuated. The added virtual slices “fill in” the missing information, by assuming a smooth transition between the adjacent sampled slices.

to calculate the surface that the centre of a virtual rolling ball traces. The idea is to roll a virtual ball with radius R touching the surface, and the regions where the ball cannot reach are considered belonging to the paper. This approach is purely discrete and requires only four scans with a small local DT mask through the volume, thus relatively cheap considering what it accomplishes. It is a kind of smoothing of the surface, where too little smoothing will let the background “dig” too deep into the paper — which is not consistent with common notions of surface, and where too much smoothing will decrease the level of detail of the surface. Finding a good compromise is necessary.

The steps needed for the 3D rolling ball algorithm are outlined below, assuming a ball of radius R .

1. Init the image with value 0 for fibre voxels, and 255 (max of unsigned byte data type) for non-fibre voxels
2. Compute the $DT(1,1,7)$ (propagates distance values away from all fibre voxels, smoothed but delocalized)
3. Mark with 0 if $DT_{image}(v) > R$ and 255 if $DT_{image}(v) \leq R$ for each voxel v
4. Compute the 3D DT (propagates towards fibre voxels to compensate for displacement)
5. Flood fill from upper and lower image boundary, and stop at voxels with

$DT_{image}(v) > R$. This will ensure that large voids inside the paper are not classified as “outside”.

Note: In principle, there is no restriction on the dimensionality of this smoothing operation. Since DTs can be calculated for data sets with higher dimensions than three, this method could be used to extract regions in N :th dimensional spaces. We would have to investigate the practical use of this.

3.7 3D pore distribution

Using a 3D DT approach, we tried to roughly label each individual pore region inside the paper, i.e., regions with no material. This would enable measurements such as pore size histogram (global size distribution), the pore size as a function of distance from the surface, etc. Note that we use *pore* and *void* interchangeably. Since pores occupy roughly 50% of the volume of paper and can be considered as the “inverse fibre network”, they influence paper properties as much as the fibres do. Similar to the individual fibre segmentation problem, pores have “wires” of fibre material and are connected to neighbouring pores. By using the smoothing approach from the 3D rolling ball method (see Section 3.6), we could simplify the problem by merging the thinnest paths and then use a 3D watershed (see Beucher (1982) for initial idea and Vincent and Soille (1991) for an efficient algorithm computing watersheds) on the DT values of the smoothed region. This will find local distance maxima and grow from these to create separate void regions. If two adjacent void regions have more than N voxels in common, these are joined into one void region (see Figure 22). A threshold $R = 7$ was determined as a compromise, to not smooth too much, to keep as many small pores as possible, and still not occupy labels for the smallest pore regions. For a small part of the labelled volume, see Figure 22.

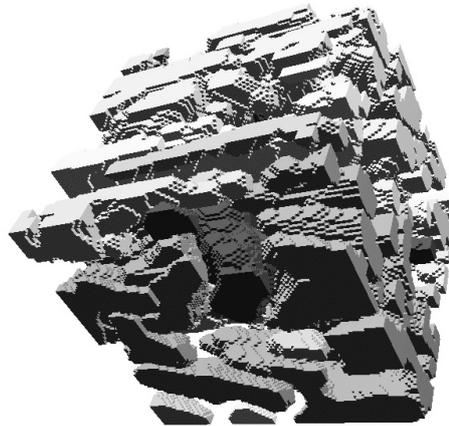


Figure 22: Pore volume rendering.

Even with the smoothing step, this pore segmentation seems as difficult as finding individual fibres. By visual inspection we realized that our watershed approach

was not able to differentiate the pore-clusters into individual pore regions with good precision. As an alternative we instead describe the pore size properties of a paper sample as a histogram of the values of the DT 3D image, see Figure 23. This is not a substitute, as we will lose the information of where the pores are distributed. It can be very difficult, even manually, to divide the pore space into individual pores. See Paper IX for details.

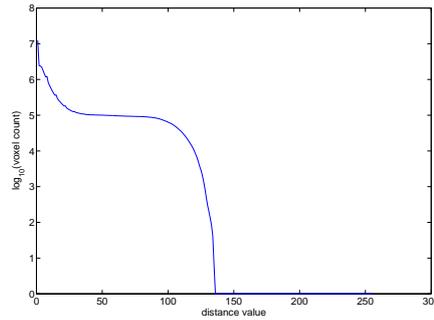


Figure 23: Histogram of DT values for the volume.

3.8 Segmenting individual fibres

As stated previously, our SEM setup combined with filtering will enable the simplest segmentation technique, i.e., a global threshold, to perform well. Transforming the grey-level image into a binarized image with only two classes, i.e., fibre voxel and void voxel, is therefore straight-forward. The next level of segmentation is nothing but trivial, segmenting individual fibres, i.e., partitioning the tightly clustered fibre material into separate fibres. We have tested three different approaches, which will be explained in greater detail in Sections 3.8.1, 3.8.2, and 3.8.3. All three approaches need user interaction, since a fully automatic approach with good robustness seems difficult to achieve.

Finding individual fibres is difficult, as each fibre in the network often have many contact areas with adjacent fibres and filler material. Therefore, using a simple 3D labelling algorithm will not identify the individual fibres. The thin long fibre structures make a 3D watershed algorithm not applicable. Instead, we need a method that is able to track the fibre path. For an isotropic 3D data set with well separated fibres that have no pores in the fibre wall, it would be a trivial task to find the winding lumen path by using any 3D fill method. This is not our case. We have tried both 2D and 3D based tracking, i.e., finding fibre contours in each slice image or analyzing the 3D image directly.

Similar segmentation problems exist in medical image analysis, e.g. separating arteries from veins, or segmenting the vascular tree in CT images. The method in Section 3.8.3, utilizes a method originally developed for medical images, but with some modification to improve the performance for fibre segmentation purposes.

As noted above, the resolution of the images is an important factor. The more anisotropic the data is, the less “3D-ish” it is. Instead, it tends to be a sequence of uncorrelated slices. However, for our data, even if the difference in resolution is a factor seven, we definitely have good correlation between slices and thus a chance to estimate corresponding points in adjacent slices.

3.8.1 Segmentation based on 2D images

The idea is to let the user mark each fibre lumen in the cross-sectional images with a mouse click. These markers will be used to automatically segment the lumen and find the surrounding fibre wall. A Graphical User Interface (GUI) was created using *guide* in MATLABTM. The more time consuming parts, “bottlenecks” of computation, are encoded in MEX-files, which is C-code with an interface that MATLAB can pass parameters through. See Figure 24 and 25 for screen dumps of the GUIs.

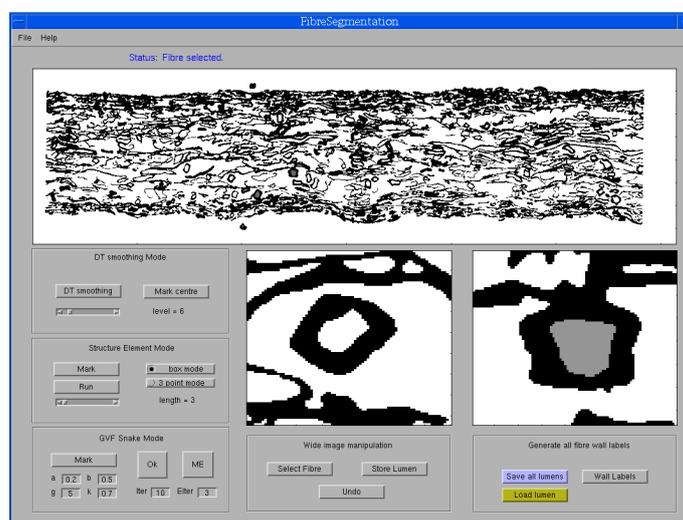


Figure 24: Interface for segmenting fibres on a cross-section.

Several 2D based approaches were used in this step. All are based on analysis of each 2D slice, one at a time, where we search for the cross-sections of the fibres. A major drawback with any sequential 2D analysis of 3D data is that we will have large problems if the fibre path tends to be parallel to the slicing direction, since we then cannot obtain any cross-sections. Since the papermaking process roughly align many of the fibres along the Machine Direction (MD), this approach is still useful for all these fibres. The fibres that are near parallel to the slicing direction will be impossible to handle. An improvement would be to combine three mutually perpendicular 2D scans of labelled fibres to obtain more of the fibre material, but a more direct 3D approach would be more robust. The large vertical forces on the paper during the press section flatten many of the fibres. Only fibres with rigid, thick walls remain basically circular. In cross-sectional images we only see a small part of the sometimes complex surface of a 3D fibre object. In combination with the natural

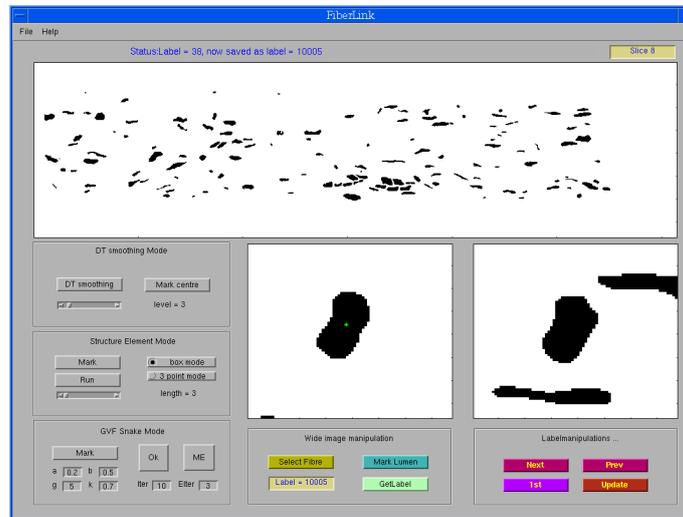


Figure 25: Interface for linking fibres from adjacent cross-sections.

shape variations which exist in many biological materials, a fibre cross-section can have quite different shapes (see Figure 26).

All our initial approaches of identifying fibres are based on finding the wall surrounding the lumen. When we continued to test this on more and more fibres, we realized that it is necessary to have a small set of different approaches. An initial approach, is to let the user decide which of methods 1–5 to be used for each fibre. This adds to the number of mouse clicks, unfortunately.

1. Well behaved fibres have no cracks in the wall, so a simple 2D fill will mark the lumen.
2. Tiny cracks can be compensated by a 2D DT smoothing (DT + threshold + DT + threshold), the same smoothing technique used to find the paper surface, but here we only need the 2D version.
3. For larger holes in the fibre walls, we use 2D snakes based on the Gradient Vector Field (GVF) extension (Xu and Prince, 1997). Since GVF is just one example of a force field that improves the convergence of a snake, we adopted this two step approach: (1) expanding forces (as in a snake-balloon-mode), (2) 2D DT gradient field, that pulls the snake towards the inner lumen wall edge. Fibre initialization was made by the user, marking one centre spot, then a set of rays are shot in 5° intervals to find the distance to the fibre wall. Due to cracks, some rays will be too long, but a median filtering of the 1D radius data will capture the main axis of orientation and we then fit a simple ellipse inside the lumen detected. Fibres with approximately ellipsoidal shape will have a good start position, but narrow eight-figure-shaped and irregular shapes with sharp corners are difficult for the snake to segment accurately.

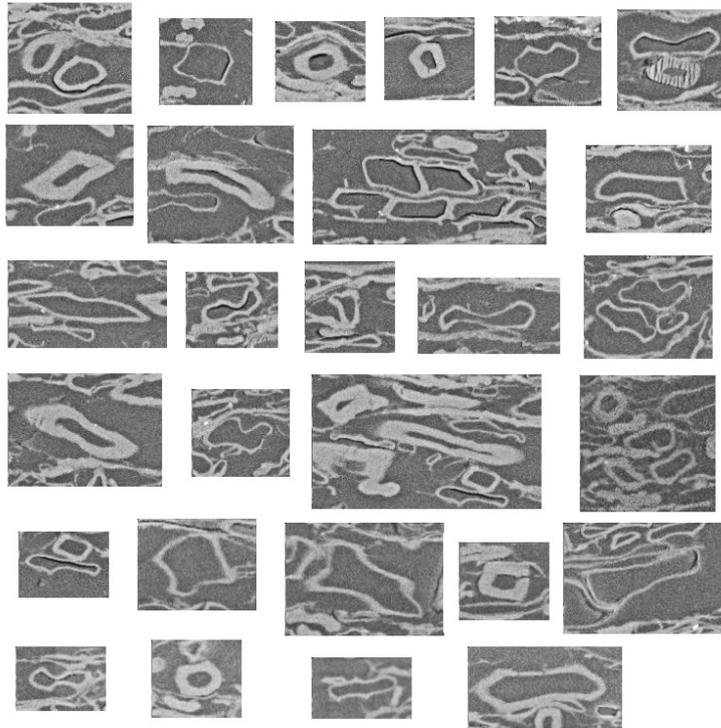


Figure 26: Examples of fibre cross-sections.

4. Oriented structure element filtering. Using thin long structure elements we could “heal” cracks in the wall, enabling the simplest fill method to perform well. By using the ellipse that was fitted to the lumen, the idea was to create a vector field that has the same orientation as the tangent vector of the ellipse at the co-ordinates on the ellipse path and with a smooth transition in-between for all other co-ordinates not on the ellipse path. Then, for each pixel the structure element that has most similar orientation to the local vector field was used for smoothing. By this technique, we could use larger oriented structure elements without smoothing so much that adjacent fibre walls mould together.
5. A few percent of the fibres were difficult for the snake to “lock-on”. For these fibres we had a “last resort” button allowing manual editing.

Note that all editing is just to guide the lumen fill method. After the fill operation is complete, the original voxels are restored. We noted that the fibres which have high eccentricity are complicated for a snake to segment well. However, the snake tests were a nice experience. They do have stability issues, but with proper setup and parameter selection, they have a way of locking on to the “target”. An interesting challenge would be to try using 3D snakes directly to find the lumen. See, e.g.,

Chapter 5 of Hamarneh (2001), where Spatio-Temporal (images with two spatial dimensions and time or three spatial dimensions) snakes have been developed. Even if they note that currently, an expert has to manually create a good initial shape, this kind of computer supported segmentation seems like a reasonable compromise between fully automatic, but less robust, and entirely manual, which often provides the highest possible segmentation quality, but unfortunately very time consuming.

Another type of 2D based fully automatic segmentation is described in Paper IV is based on DT histogram and some basic features for finding lumen candidates. An example of a successful result can be found in Figure 27. The variation in segmentation quality varies quite a bit, making it less ideal for use on the entire 3D image.

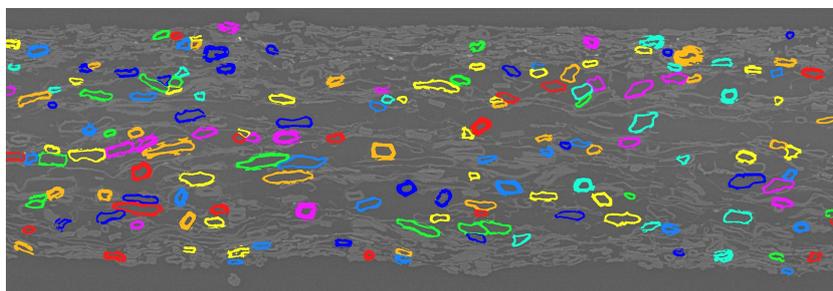


Figure 27: Example of fibre segmentation result.

The oriented structure element approach gave mixed results. For some fibres they work perfectly, removing the small gaps in the fibre wall, but for others the shape was altered too much. Both this and the snake approaches have a set of parameters that need individual settings for *each* fibre. The obvious drawback of this approach was the extremely tedious manual editing, resulting in sore eyes and headache. There always seemed to be some 2D piece of each fibre missing, resulting in problems for the 3D reconstruction, since we needed a way to take care of the missing parts. In conclusion, all five methods are not feasible for such a large data set as we have. After working with the different 2D techniques, we were even more convinced that anything else than true 3D methods are not up to this task. If the 2D approach would have worked well, it had been possible to track fibres using a less powerful computer than 3D based methods require.

3.8.2 3D vector path estimation based tracking

The “debacle” above was forgotten and it was time for new and improved ideas, in three spatial dimensions of course. The idea is to still take advantage of an expert operator to identify fibre lumens in a cross-section in the first z -slice, and then the method should try to track this fibre lumen through the volume in a fully automatic way. Remember that we define the z -axis to be along the direction of microtome slicing, not perpendicular to the surface, which is common.

Our implementation does the following: 3D DT on the background voxels. The

wall voxels v have $DT_{image}(v) = 0$ and the lumen and background voxels will have increasing values the further away from any wall voxel. This DT_{image} space is now used to guide the “probe” we send into the lumen of each fibre. The initial direction estimation is perpendicular to the cross-section, $v_i = [0, 0, 1]$ and this is updated along the path. In each step we allow only steps in \mathbb{Z}^3 inside the local $3 \times 3 \times 3$ neighbourhood and also less than 45° from the current estimated direction vector v_i . This enables a fast table look-up in the local neighbourhood of 26 voxels, to obtain a list of the voxels which have the highest DT_{image} value (may be more than one) and select one among these. We want to stay away from the walls which seem sound, to “keep in the middle” of the lumen as we trace it. Many of the fibres are elliptically shaped, hence having lumens with high eccentricity. This cause problems, since the DT only calculates the distance to the closest wall voxel in each position. We get good “steering” along the minor ellipse direction, but rather loose steering along the major ellipse direction, producing a wiggly path inside the lumen. See Figure 28. To improve the performance, we added the constraint that we may not deviate further than M voxel units away from a line segment, fitted to the N last voxels of the path. During the experiments these were manually chosen, as e.g., $M = 5$ and $N = 5$. Our problem was that the direction estimates were not consistently accurate. Our restriction was then of little use and the path often diverged into the paper void through a wall crack or one of the artificial holes caused by the low sampling along the z -axis. It is however, a true 3D method, which could handle fibres with any path through three spatial dimensions, as long as we do not loose the tracking, which happened rather often. If an improvement of the local fibre path estimation is possible, maybe this idea could be developed into something really useful.

3.8.3 Ordered region growing with a fibre touch

Ordered Region Growing (ORG) has been used with great success on medical images to segment veins and arteries (Adams and Bischof, 1994; Yim et al., 2000). We have tried this method with the only modification that the steering function f is adjusted to suit our problem domain. For the rest, the method was unchanged. It was encouraging, to find an existing method that was well suited for our problem! ORG was designed to segment tree-like structures as e.g., blood vessels. An expert operator manually seeds one start-voxel and one or several end-voxels, for a specific part of a vessel tree. Then ORG finds the path in-between by incrementally adding voxels that are optimal, based on the steering function f . Since paper fibres do not branch, we only need a single path, which is a special case of a tree, so it is not a problem.

The major difference to our method based on estimating local orientation to track fibres, is that ORG can recover if it gets lost, since it continuously checks which voxel to add by evaluating f on each boundary voxel of the current tree. The speed and accuracy of the method depends only on one thing: selecting a good function f , which is used to decide which voxel in the boundary set we will use in each iteration. A good f will grow the path along the desired fibre lumen, and not allow the tree to deviate into other regions.

The ORG finds a path through the fibre lumen, but cannot segment either the entire lumen, or the fibre wall. Therefore we need an additional method to handle

that. In Paper III we have developed a technique to segment 2D fibre walls, which works reasonably well. We extended that idea to three spatial dimensions, to get full fibre wall segmentation. The details follow after the description of how the ORG method is implemented.

Finding a path through the lumen of a fibre

To initialize the computations for each fibre, ORG requires the user to manually mark two voxels in the volume, the preferred start- and end-voxels that must be inside the lumen of the fibre. The conclusion in Paper III that lumen tracking is easier and more robust than wall tracking, still holds for images with three spatial dimensions. It is not straight-forward to implement robust automatic seeding, instead we have concentrated our efforts to adapt ORG to our fibre tracking purposes. Without having tested the following somewhat vague idea, a kind of seeding with a skeleton structure of the void regions may provide an opportunity for automatic seeding. However, to be robust, this will need to be fine-tuned for this specific seeding problem.

Ideally, we only want the optimal path from start node to end node, but depending on how well it is possible to design the steering function f , which of course depends on how difficult the data set is, we obtain a tree structure with more or less branches. Without a well designed f , the method is “lost in the void space”. With a reasonable f , the ORG can create a small tree structure and as soon as the end node is included in the tree it stops. Each node has only one parent node, so we can back-track from the end node along a unique path to the start node.

Medical images often have vessel trees with high grey-levels in the middle and decreasing towards the walls of the vessels. This is achieved by injecting a contrast agent into the patient before imaging. In this case, the choice of a steering function $f(\mathbf{v}) = \text{greylevel}(\mathbf{v})$, can be a good initial approach. Our data are nearly binary to start with, so grey-levels are not optimal for f . Instead we use the following: By a DT we get the distance from the closest non-void voxel for all void voxels. We want the path to stay away from the walls, favouring a path centred in the fibre lumen. Roughly elliptical cross-sections will have a set of local DT_{image} maxima inside the elongated lumen (see Figure 28), so the path may oscillate during tracking. This is not a severe problem, since we can attenuate these oscillations in a post-processing step, by using approximating splines. Another crucial modification to f was that we want to favour local steps toward the end-voxel. Combining the need to favour voxels centred within the lumen and steps towards the end-voxel, we found the formulation of f as in (1) useful:

$$f(\mathbf{v}) = DT(\mathbf{v}) + \frac{1}{\|\mathbf{v}_{end} - \mathbf{v}_{start}\|} (\mathbf{v} - \mathbf{v}_{start})(\mathbf{v}_{end} - \mathbf{v}_{start}) \quad (1)$$

Note: The parameters of the Chamfer DT will affect the selection. For example, using larger chamfer weights will favour steps centred inside the lumen, compared to steps towards the end. This is not a problem, since a simple scaling of the DT values will compensate for the change, but is worth keep in mind.

Improvements could be to allow more difficult cases by marking a *sequence* of voxels, then run ORG iteratively on each subsequent pair. This stricter guidance

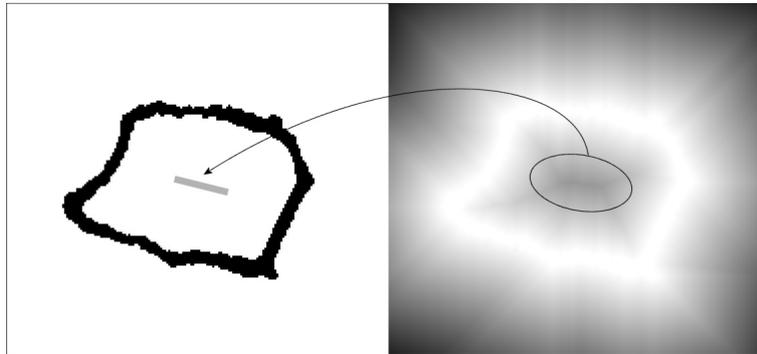


Figure 28: Illustration of how the ORG path tracks in a 2D cross-section that is approximately perpendicular to the fibre path in three spatial dimensions. (left) Sample fibre, (right) DT_{image} of fibre, where higher DT_{image} values are darker and lower DT_{image} value brighter. The DT_{image} maxima are located along a line in the right image, and a corresponding grey line in the left image show where the ORG path is most probable to scan.

will decrease the risk of ORG loosing track and getting lost in the void between fibres, but this mode requires more user interaction. Some method details follows. During the calculations, three sets R, G, B - keep track of the progress, where R is the region of all voxels we have visited, G is the growth, i.e. voxels added in the current iteration, and B contains the boundary voxels, i.e., the subset of voxels in R that have a at least one neighbour voxel which is not in R .

- Setup: Insert start node in R and B , then in each iteration:
- Select a new seed node s_n , from B , such that s_n have the highest $f(s_n)$ of all voxels in B . Using a poor decision function f , has the effect of increasing the size of especially R , and B , since we step in many directions that are not optimal paths towards the end-voxel.
- Each neighbour voxel of s_n will have s_n as its parent. This simple update, have an important property that throughout the entire region R , every voxel will have a unique parent. When tracing backwards from a voxel in R , there will be only one path to the start node.
- Update the three sets:

$$\begin{aligned}
 G_n &= \text{Neighbours}(s_n) \setminus s_n \\
 R_{n+1} &= R_n \cup G_n \\
 B_{n+1} &= (B_n \cup G_n) \setminus s_n
 \end{aligned}$$

The iteration continues until the end-voxel is reached. After that we trace back, following the parent pointer in each node, and save this path in a list of (x, y, z)

co-ordinates. To efficiently find the voxel in B with highest $f(s_n)$, we keep the set B sorted according to f . Since the set is internally stored as a linked list, we can simply retrieve the first element in each iteration, a quick operation. To minimize the space complexity, we encode the parent as a lookup into a 26 item table since we always have 26 neighbours and we can say, e.g., the local step $[-1, -1, -1]$ is encoded as 1, thus only five bits are needed (four bits can only encode $2^4 = 16$ combinations, which is too few). Each (x, y, z) position is stored as the memory-offset. We use $\text{offset} = x + \text{xsize} \cdot (y + \text{ysize} \cdot z)$. Using four bytes allow us to handle volumes with up to 4 G voxels.

Lumen segmentation from the ORG path

From the ORG, we now have a somewhat wiggly path inside the fibre lumen. The idea is to use this as a base for finding all lumen voxels. For this step we use the 3D SeparaSeed method (Tizon and Smedby, 2002). A probably more known method that performs a similar kind of segmentation is the *Watershed* based method (Beucher, 1982; Vincent and Soille, 1991). Contrary to Watershed, in SeparaSeed the pixels do not need to be sorted and only a small local mask can be used to iteratively update the connectivity. Thus, it is often much faster than Watershed and well suited for our purposes. Each labelled voxel has a connectivity value that is a measure of how “close” it is to the seed region with the same label. Since we use two classes of seed-points, void and lumen, we will classify all background voxels according to if they belong to lumen or background. All wall voxels are considered a restricted area, so no paths can pass through these. As connectivity between two voxels, we use a number based on a DT from the non-background (fibre) regions. Cracks and holes in the fibre wall that are smaller than the diameter of the lumen will not let the void region flow into the lumen, as the connectivity is larger to the lumen seed voxels. See Figure 29 for a synthetic test case with a 2D fibre with some holes in the fibre wall has been labelled using 2D SeparaSeed. It is easier to show a 2D test case to view the basic principle, and the 3D SeparaSeed works in exactly the same way. That the extension from 2D to 3D is straightforward for SeparaSeed, is an additional good property.

Connectivity between voxels A and B is defined as: among all possible paths from A to B , the path with the highest strength is used. Strength is defined as the lowest DT_{image} value in that path. With SeparaSeed, the optimization problem can be efficiently handled using chamfering, which is not immediately obvious. To be able to measure the connectivity to void voxels, we need to find a set of them to use as seed voxels and this is done by the following procedure: a DT from the ORG path is thresholded between two radius parameters, R_1 and R_2 . The non-wall voxels inside this region are considered to be void voxels. This is based on the assumption that fibres are approximately circular and with diameter less than $2R_2$. Since we use the ORG path, this shell of surrounding voxels will follow the fibre, and even if the path is a bit wiggly, the DT will have a smoothing effect, so the shell is rather smooth. Now we have the seed regions for the void and the lumen, so SeparaSeed can find all lumen voxels. Problematic cases are fibres with thin lumen, since then the connectivity measure is not suitable.



Figure 29: An example of 2D connectivity labelling using SeparaSeed. (left) original fibre with ORG path in lumen, (middle) the surrounding DT_{image} -shell created from the ORG path, (right) result of SeparaSeed – lumen is now separated from the background.

Wall segmentation from the lumen region

From the lumen voxel region, we need to mark the surrounding wall voxels which belong to the current fibre. Any adjacent fibre wall voxels should be left out. Here we extend the idea, presented in Paper III, of using a histogram based approach based on a DT from the lumen region. A histogram for the ratio between fibre wall voxels and the total amount of voxels in each DT_{image} layer is calculated. Since the total number of voxels are different for each DT_{image} layer, it is not a good idea to only count the number of wall voxels in each DT_{image} layer. The DT_{image} layers for the $DT(3,4,5)$ are defined (Borgefors, 1994) as $\{(0), (3), (4,5,6), (7,8,9), (10,11,12), \dots\}$, i.e. shells around the lumen region. The assumption is that the fibre wall is approximately of constant thickness. If so, this histogram will have a step-edge shape, see Figure 30, and the average wall-thickness can be estimated. To allow for some local thickness variation, we have a parameter *scale*, that allows to fill up to $scale \times wall\ thickness$ from the lumen. We now want to use our calculated local wall thickness. We use a restricted reverse labelled DT, where only fibre-material voxels will be used in the propagation of labels. It is initialized by setting the lumen voxels to $scale \times wall\ thickness$. The reverse DT will then fill the fibre walls, by locally decreasing the initial DT_{image} values for all of the voxels, and at the same time, set the voxel label.

We thus use a three-step method to find a fibre: (a) find a set of lumen voxels that trace the fibre through 3D space; (b) use the path in (a) as seed voxels to find the entire lumen; (c) from the lumen found in (b), we can find the surrounding wall voxels. The result of the method is a voxel volume where for each $fibre(i)$ we found, two sets of voxels are associated: the set of lumen voxels with label $2i$, and the set of wall voxels with label $2i+1$. This labelled volume is then used as the base for the successive measurements. By this approach we could find 25 of the cylinder shaped fibres in the dataset. This is about 10–20% of the total amount of fibres

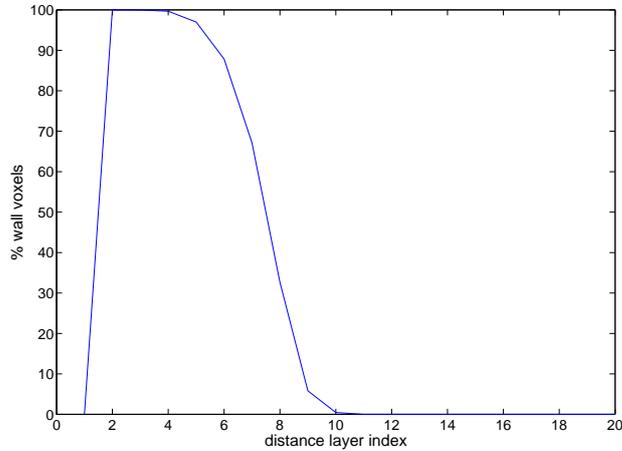


Figure 30: A DT_{image} layer histogram based on a 3D image, displaying the percentage of wall voxels in each distance layer surrounding the set of segmented lumen voxels. See the fibre in Fig. 31.

present. Some of the remaining fibres are very short, these we have not segmented, and the other fibres have more complicated shape, and are thus difficult to segment accurately. An example of a segmented fibre, can be found in Figure 31, and in Figure 32 six more are displayed.

3.9 Measurements

After the digitization of a paper sample into a digital model, the actual use of the model can start. For ease of computation, we resampled to cubic voxels for the measurements. We have also developed two representations to aid in this measuring step, since they enable more efficient measurements than if we use the fibre model of voxels directly.

Our methods can provide alternative methods of measuring paper and fibre properties, with some measures that are possible with other techniques, and some that are unique to our approach. There are most certainly other measurements possible than the ones we provide here, but as a sample collection, they provide an indication of what can be measured. Especially the extreme level of detail for each fibre is unique. Even considering that all measurements will be biased by various errors from the filtering, registration, and segmentation, there remains a good deal of information that should be useful.

We define three levels of measurements, since each have different requirements and produce estimates on different scales. They are *paper level*, *fibre network level*, and *fibre level*. On the *paper level*, we have three measures: surface estimation,

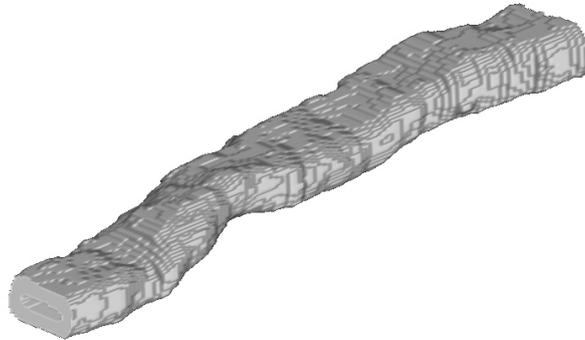


Figure 31: Example of labelled 3D fibre, found by the methods in Section 3.8.3. This fibre has a rather thick wall, but still the fibre bends and twists a bit along its path.

paper density, and pore size analysis. Neither of these requires individually labelled fibres, only a binarized digital 3D image of a paper sample. We have $0.7 \mu\text{m}$ resolution in the xy -plane, a detailed study of the surface is possible, even if some smoothing during the 3D rolling ball step will occur. Density is possible to measure as soon as the surface is determined, since we already know which voxels are void and which are fibre material. The 3D void analysis enables an estimation of the void regions inside the paper sample, which we use to make histograms of distributions along the x , y , and z axis, respectively.

At the *fibre network level*, we have implemented code to count the number of contact points between fibres, measure free fibre length, and contact area. This gives some information about how the fibres are interconnected.

Finally, we have the *fibre level*. Most of our measurements have focused on this level, since the individual fibres are our main concern. It is possible to estimate at present: length, wall thickness, curvature, twist, torsional rigidity, aspect ratio, degree of collapse, lumen volume, fibre wall volume, and bending resistance. With this set of measurements, we can capture some basic fibre properties at a very high level of detail. The surface representation of the fibre wall is ideally suited for, e.g., wall thickness variation analysis. Since it also captures the shape of the fibre wall, as well as the local thickness, several other measures could benefit from this. The curve representation of the fibre lumen is a one-dimensional curve through 3D space centred within the fibre lumen. Thus it is great for fibre length measurements and for keeping track of where through 3D space the fibre path is traversing. Combining

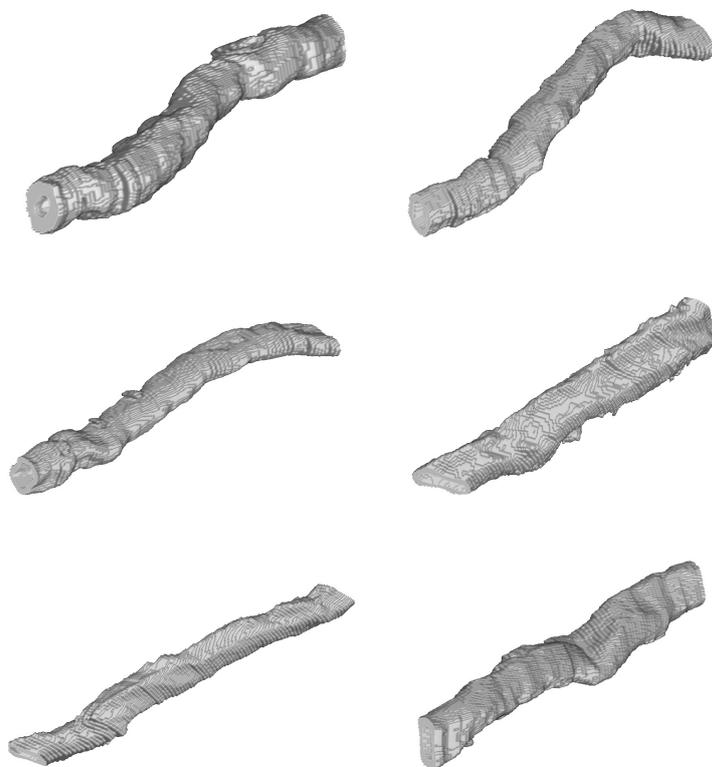


Figure 32: Example of labelled 3D fibres. Both the lumen and wall are identified, but here we only display the wall voxels.

these two representations with the completely labelled 3D image should provide further possibilities of measurements on both the fibre and network level. The size of our sample is far too small to analyze larger level effects as e.g., flocculation. We need either to digitize a larger piece of paper or restrict the analysis to only small sampled regions. As there are other techniques better adapted to the scale of large level effects, we feel our approach is not useful for that kind of measurements. See Hansson and Johansson (1999) for an image analysis based approach to measure surface topography of paper samples of size $5 \times 5 \text{ mm}^2$, which seems extendable to larger samples. For details about our measurements, see Papers VI, VII and IX.

3.10 Rendering 3D images

An important aspect when analyzing data is how we select to view the results of a computation. For one-dimensional scalar numbers, a colour diagram can be much more useful than a simple dump of the numbers to a table. If the result is a two-dimensional array of scalar numbers, we can use a 2D diagram with bars, or map

the 2D array directly to the screen as an image, e.g., using blue for low valued numbers and red for high valued numbers, resembling a temperature scale. This is intuitive and easy to understand, which makes the interpretation of the results more straight-forward. When the results of calculations has higher dimension than two, the presentation is more difficult. The basic restriction is that our visual system is two-dimensional. If we have three spatial dimensions or even higher dimensional data, we need to project the data onto a 2D display. Hence, we will loose some information, and it is necessary to decide what can be left out. See Computer Graphics (2002), for a list of many research groups that are working with computer graphics.

For image analysis of 3D data, rendering is an essential part, since we need to “look at” our results, see Figure 33. A simple alternative can be to browse through the volume, looking at one 2D slice at a time. However, doing so much of the overall 3D information is lost. Computer graphics is a huge and active research area, even only looking at the part that deals with scientific visualization, i.e. presenting results as visual information. Even if it is a bit overwhelming, it is also great fun. We have taken the opportunity to see how computer graphics can aid us in visualizing 3D fibre images.

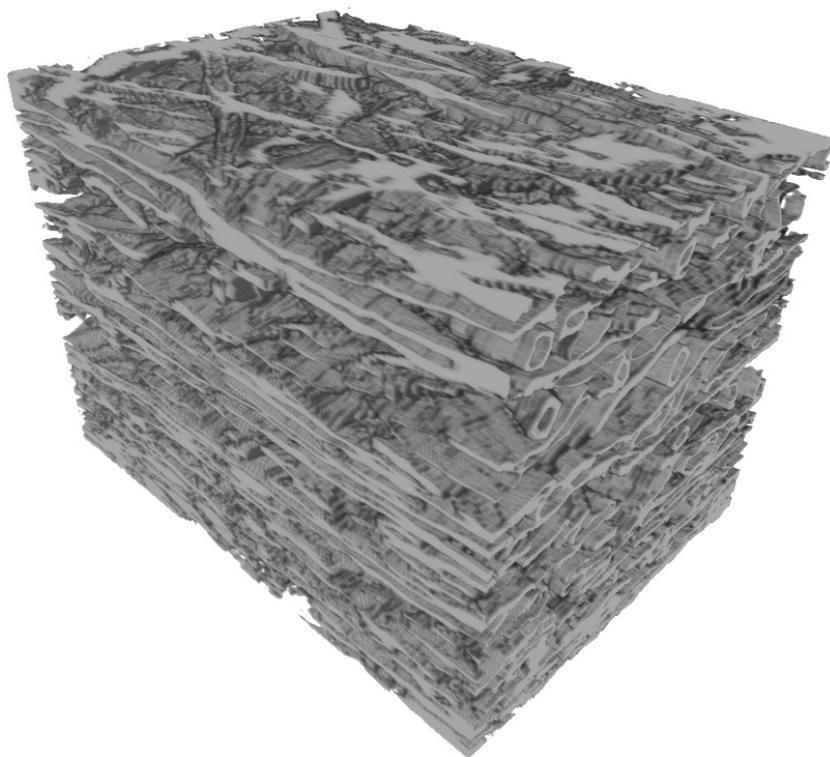


Figure 33: An interior region of the digital voxel model of the paper sample, size $500 \times 500 \times 102$ voxels.

The first classification of rendering methods is probably into volume rendering

or surface rendering. If your objects are opaque and you do not care about their interior, it is often possible to have a compact representation by only storing the surface elements. For surface rendering, you may choose between voxels or patches of, e.g., triangles. Rendering 3D volumes often involves moving huge amounts of data between a hard disc and the RAM. This severely limits the performance of the rendering, so compression techniques and efficient data structures are important topics. Even simple compression techniques, e.g., Run Length Encoding (RLE), can make a big difference. There will always be volumes that cannot be loaded directly into RAM. That is why many develop techniques for what is known as “out of core rendering” to enable rendering of volumes that fit on your hard disc but not in your RAM. Multiprocessor systems are ideally suited for the rendering task, albeit algorithmic design becomes even harder. An indication of how large models researchers are interested in can be found at the web site for the “Visible Human” project at NLM (2002), where 15 Gbyte of cross-sectional images have been assembled into a volume model. Another impressive site to look at is “The Digital Michelangelo Project” (Levoy, 2002) where the David statue in Florence was digitized into a stunningly large set of 2,000,000,000 surface polygons! Without the continuing development that improves current techniques, these data sets would not be possible to visualize.

Some of the rendering methods described below need triangular patches as input. We use the Marching Cubes (MC) method to create a triangulated surface for our voxel data (Lorensen and Cline, 1987; Wyvill et al., 1986). There are other alternatives (see, e.g., the Tetra-Cube method in Carneiro et al. (1996)), which needs to be combined with MC since it alone creates even more triangular patches. There are techniques to post-process the triangular mesh, to reduce the huge number of triangles (triangle decimation) without changing the geometry too much. We have not tried such techniques. Another important distinction of methods is the relative size of the voxels compared to screen pixels. As the voxel models are getting larger, it seems pointless to run MC and get up to four triangles per voxel, when the voxel size projected on to the screen is smaller than a single pixel. Because of this, the idea of “splating” was invented, treating each voxel as a point which is splatted on to the screen with a certain 2D foot print (Mueller and Yagel, 1996). The trick is to select a foot-print that does a small amount of smoothing and simultaneously avoid creating artificial holes in the rendering of opaque objects.

Below is a list of rendering methods and storage schemes that we have implemented code for. POV-Ray and QSplat needs converters: from voxel to POV-Ray format, and voxel to PLY format (QSplat supplied a converter from the PLY polygon format to the QS bounding spheres format). The ShearWarp method is implemented from scratch using C++/OpenGL, although somewhat annoyingly, there are some “bugs” left to kill until it becomes useful. Code for converting between the shell format and 3D rectangular grid volumes has been implemented, although the final rendering of shell voxels is not finished. Finally, thanks to a demo-version of VolView we can actually present some good-looking renderings of our paper volume and fibres ...

POV-Ray™ (Persistence of Vision Ray Tracer) A freeware ray tracer (POV, 2002) that many use for creating computer art, see www.irtc.org for some impressive ren-

derings. It is not ideal for visualizing large volumes, since both the ASCII format of the POV-Ray files and the computationally demanding rendering approach will take a lot of time to create a single image. Ray tracing is definitely aimed at non-interactive rendering, that is, you start a rendering, go to sleep, and in the morning your rendering is finished. Nonetheless, it is a fun program to experiment with. The small fibre in Figure 21, was created by using shape-based interpolation and POV-Ray. The volume was approximately 100,000 voxel large, but the POV-Ray file created was 16 Mbyte!

Shear Warp A clever way of performing volume rendering without any specialized hardware. Even if it is now commonplace to have highly accelerated 3D graphics card in a standard PC, these are optimized to render large amounts of texturized triangles, not fast rendering of voxels. This would most certainly change if, e.g. Quake IV would require voxel rendering. The basic idea of shear warp is to factor the projection from three spatial dimensions to two spatial dimensions into two different steps, a simple shear that is optimal considering how the data is stored for a volume, and the warping that can be done by OpenGL hardware, together forming a very fast method. Both orthographic and perspective mode is possible, although perspective mode takes more time, as then each 2D slice need to be rescaled in the shear step. Shear Warp has even been implemented in hardware to provide real-time volume rendering of 30 frames per second (fps) on a standard PC, for volumes up to size $512 \times 512 \times 512$ voxels. See AVL (2002) for the graphics card and Lacroute and Levoy (1994) for details about Shear Warp. Our implementation of Shear Warp is in a state of massive bug-removal at the moment, but this method is definitely our favourite approach to voxel rendering. Especially for volumes where the surface-to-volume voxel ratio is very high.

QSplat This technique of Rusinkiewicz and Levoy (2000) is a variation of “splatting” where each voxel centre is shot at the screen and “splats” into a small footprint. This ensures that we have no holes in the rendered image and also enables an efficient traversal of the data. QSplat improves this further by adding a clever “Bounding Spheres” data type, that is a kind of hierarchical structure of different resolutions that enables a successively improved rendering, while maintaining interactivity. Note that splatting may produce extremely bad images when zooming close into the model, since the method consider a voxel as a single flat “blob”, not as a well defined 3D object. A solid implementation can be downloaded from Rusinkiewicz and Levoy (2000). It is fast, multi-threaded (the GUI can handle user interaction independently of background calculations) and they have a few QSplat models to download and test the renderer on. QSplat is not limited to surface patches, but in the downloadable program they use a triangular input using the PLY polygon file format, since this simplifies the estimation of normals. To be able to use this software for our fibre volume, we wrote a small conversion utility.

Shell data types As many 3D voxel objects have rather opaque surface voxels, it makes sense to only store these, and not all the internal voxels as well. Even for semi-transparent objects with a non-binary opacity function, the innermost parts may be invisible or contribute very little to the rendered image, thus unnecessary to store. A problem with abandoning the 3D rectangular grid storage scheme is that we need to store the (x, y, z) co-ordinates explicitly for each voxel. Without an

efficient storage scheme, the size of the volume may increase instead of decrease, even if we take into account that the surface voxels are often less than the total number of voxels. The trivial data type for storing a shell, is a 1D list of quad-tuples, containing (x, y, z, value) . As an example we use a 8-bit grey-scale volume with $1000 \times 1000 \times 1000$ voxels, where 1% of the voxels are surface voxels, i.e., $10E6$. Using ordinary implicit co-ordinates and storing all voxels, we need 1E9 bytes. If we use the “trivial” shell data type we would need: $10E6 \cdot 7 = 70E6$ bytes, thus a compression about a factor 14. We cannot get more compact storage with this data type, as each quad-tuple (x, y, z, value) needs $3 \cdot 2 + 17$ bytes. A more compact storage of the surface voxels is achieved by grouping the co-ordinates in a way that minimizes repetitions, and adding a small 1D or 2D header of pointers needed for the co-ordinates we have removed. In *Slice-based Binary Shell (SBS)* by Kim et al. (1999), two 1D arrays are used, one for z -index and the other for a list of (x, y, value) triple-tuples of data. A random lookup of voxel (x, y, z) will proceed as follows: check z -index for start- and end-index for all voxels in the slice, which is two quick lookups: $z\text{-index}[z]$, $z\text{-index}[z + 1]$. Since the voxels in that z -slice are sorted according to $x + x\text{size} \cdot y$, we can do a binary search of (x, y) , a $O(\log n)$ operation. Continuing the example above, we may now store the volume in: $10E6 \cdot 5 + 1000 \cdot 4 = 50,004,000$ byte, a compression about 20 times. The 1D array of z -index, is used to aggregate all voxels with a common z value, in the 1D data array. This can be taken one step further, by using one 2D array of pointers and store the surface voxels in a 1D array of pair-tuples (x, value) . Then we would need: $10E6 \cdot 3 + 1000 \cdot 1000 \cdot 4 = 34,000,000$ bytes, a compression about 29 times.

In Udupa and Odhner (1993) and Grevera et al. (2000), more information is added to the (x, value) data tuple, to have information about the neighbouring voxels opacity available. This lower the compression possible, but speed-up the rendering as this extra information can be used to decide if the neighbours hide the current voxel. Note that if we had chosen, e.g., 14% surface voxels in our example, the space needed for the most compact shell data type would only compress a factor two, compared to the implicit co-ordinate volume. Thus, how much it is possible to gain by using a shell data type depends on the surface-to-volume voxel ratio.

VolView A volume visualizer made by KitWare, who also maintains the open source Visualization Tool Kit (VTK), a popular graphics language on a higher-level than, e.g., OpenGL. The latest version 1.3 of VolView is now very fast. VolView can use both the hardware accelerated texturing features of a modern graphics card, and supports the extremely fast voxel accelerated VolumePro graphics card. The rendering in Figure 34 uses half of our dataset (115 Mbyte) and takes about 20 seconds on a 1.3 GHz Athlon PC. Noncubic voxels can be rendered, although we had to set a voxel size of $1 \times 1 \times 5 \mu m^3$, instead of $0.7 \times 0.7 \times 5 \mu m^3$, as otherwise the renderings looked more like stacks of slices than a volume. The cause may be a normal estimation that cannot handle our anisotropic data directly.

The conclusions of our experiments are that digitized paper has such a high surface-to-volume voxel ratio, that it is not a good idea to save the surface in a shell data type, as that will at best compress about a factor two, but make random voxel access slower. The low resolution in the z direction increases the surface voxel count, as the larger movements will unveil more fibre material voxels to the back-

ground, in combination with the fact that the fibre walls are so thin. Approximately 14% of all voxels are surface voxels in the paper model. In applications where more solid and thick structures are to be stored, may have only a small fraction as surface voxels, thus an opportunity for compression exist. We believe that volume rendering using ShearWarp or QSplat is well suited to our images.

To create good-looking renderings, accurate surface normal estimations are essential. The basic technique is to run a central difference operator on the voxel data. While this is trivial to implement, the result may not be as good. More advanced methods use a larger local neighbourhood to estimate the normal direction for each surface voxel. Estimating surface normals can be rather time consuming, so if we need to store these vectors as well as the fibre labels (to lower the rendering time), even larger data structures are needed.

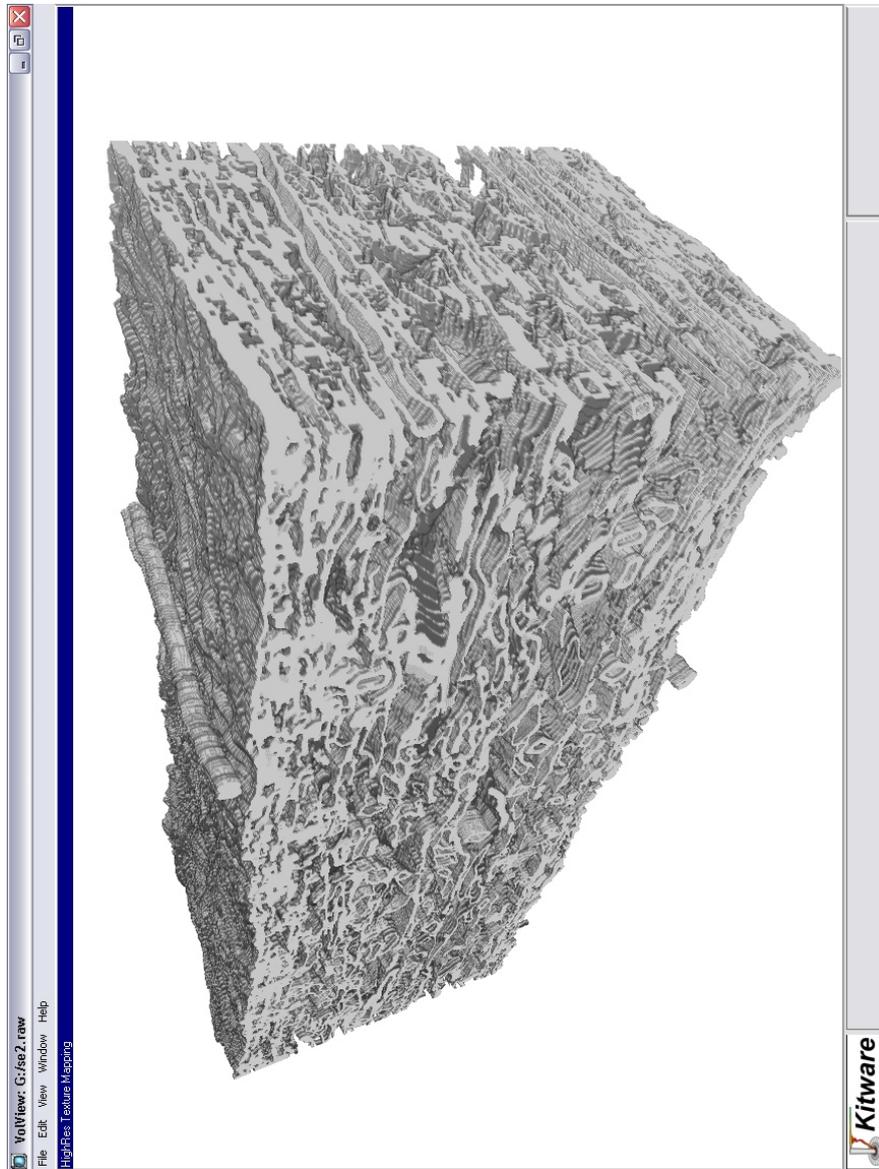


Figure 34: One half of the digitized paper, size $1500 \times 768 \times 102$ voxels. The upper rayon fibre is visible, and some of the internal fibres have been split by our virtual cutting, revealing the inside of the fibre wall.

4 Summary of the included papers

For those that prefer to have a slightly more detailed alternative to abstracts, but shorter than the full paper, this section will give you the main content of each of the enclosed papers.

4.1 Paper I: Comparison of two different approaches for paper volume assembly

The optical light microscopy data set from STFI and the SEM data set from Stora-Enso were compared, to evaluate which set is the most appropriate to use for individual fibre segmentation of 3D images. Since both approaches require microtomy, i.e. slicing of thin sections iteratively to unveil the inner structures for the 2D image capturing device (CCD camera or SEM electron counter), similar errors caused by the cutting are introduced in both methods. Thus we compare the imaging techniques only. Due to the high demands on low geometrical distortions, we prefer the SEM setup, as it is then possible to image each 2D image slice at the remaining paper block of the sample, which will result in much less geometrical distortions of the sample. During the microtomy process both the paper block and the sliced off thin section will be distorted by the forces from the knife. Since the surface of the paperblock has a large support, it can recover from much of the elastical deformations, while using the cut-off thin sections directly would be a less ideal choice for imaging. By small distortions we mean below 2–3 μm in our case. Even if the CCD camera has higher contrast, it was of no advantage to us. On the contrary, it makes the segmentation between void areas and fibre material much more complicated due to less bi-modal histogram shape. We conclude that the SEM setup is superior and we will use this in our further studies.

4.2 Paper II: Minimizing scanning electron microscope artefacts by filter design

As noted in Paper I, the contrast of SEM images are quite low. This is especially true as we have a time constraint due to the fact that we need to capture a large number of 2D images for the 3D image reconstruction. That is why we experimented with a new type of filter that performs a non-linear filtering of the image to both preserve edges and smooth the image into a near binary image. There already exist many approaches to this problem, but nonetheless, this new filtering approach can be used as a complementary technique for blurred binary images. Compared to other filters, e.g., median image filtering, this filter is faster and has the same straightforward implementation. The idea is based on the following: we assume the image is a binary image which has been blurred by a point spread function (PSF), caused by combined camera filters during the imaging. Making an initial segmentation by hand, we select the fibre objects and perform a distance transformation from these pixels. Using these DT_{image} values, we can obtain an average profile of how a step transition (the assumption) has been blurred by the PSF. The second part is to use this estimated PSF to restore as much as possible of the original unblurred image. Taking advantage of our DT approach we do not calculate the inverse PSF and

convolve the image by a discrete mask. Instead, we can reuse the DT_{image} values and the DT_{image} profile. Simple lookup and subtraction of the profile will ensure that, on average, the adjusted DT_{image} profile will be a step function. We tested and compared with different convolution based approaches, as to optimize a discrete kernel mask, Wiener filtering, and iterative blind deconvolution. The developed method is simple to implement, has only one parameter (the initial segmentation) and was able to attenuate the background variance in our images by 50%.

4.3 Paper III: Slice based digital volume assembly of a small paper sample

Having a stack of unregistered but filtered 2D images, we would like to assemble them into a registered 3D image. Since there exists no ideal devices for capturing 3D images of paper, it needs quite a long procedure, since we first have to split the paper sample in parts during the cutting process and later reassemble the 2D images into one 3D image.

As the paper material is often a thin sheet, it is preferable to create 2D cross-sectional images that have a higher aspect-ratio than the standard CCD resolution of 1024×768 . By taking three overlapping images at each cut-level, we can register these into one 3072×768 image. To get a reasonable fast and accurate registration we use Hierarchical Chamfer Matching, with a final tuning and comparison done with local correlation. The next registration step is more difficult, since the relatively thick slices we cut off will make adjacent images quite different. Using correlation or template matching is not robust enough for this. Instead, we used landmark based registration, which assumes knowledge of some nodes in each 2D image that should be connected through the 3D image as a smooth path. Since a fibre can bend quite much, and may not exist in all paper slices, we manually selected a few “good” ones. We used the two straight Rayon threads that were positioned at the paper surface explicitly for registration purposes. Restricting the transformation allowed to be affine, i.e. a linear transformation on the form $[u; v] = [a, b, c; d, e, f] \cdot [x; y; 1]$, we were able to obtain an accurate registration even with rather few landmarks. Evaluation was done by visual inspection. We only detected a few areas which had local movements above 2–3 pixels. The remaining, i.e., most, regions of fibres had a smooth path through the 3D image.

4.4 Paper IV: 2D Segmentation and labelling of clustered ring-shaped objects

The methods described in Papers I, II, and III, have reconstructed the 3D structure of paper. Hence, we now have a voxel based representation of the paper sample. From this we can have the information each voxel, if it is fibre material or void. From this low level information we want to detect as many fibres as possible, by clustering together sets of voxels that belong to a specific fibre. This is a complex task, since the shape variations for fibres are quite large. After some problems, we realized that it is less difficult to create a fibre detector if we concentrate on the fibre lumen instead of on the fibre wall. Even if this approach will exclude fibres

with no lumen (no interior void), we believe it is a key to good segmentation. To further simplify the segmentation we analyse a 2D image, i.e. cross-sections of the 3D image. Working on the 3D image is preferable, but much more complicated. The basic idea is to segment a set of lumen candidates. For each of these we try to assess if it has a surrounding wall of fibre material. If so, we assume it is a fibre. By calculating a local DT around each lumen candidate and measure the ratio of void to fibre material pixels at each distance value, we obtain a profile of the fibre material density. For an ideal case, a fibre with constant wall thickness, no adjacent fibre walls and no cracks (open pores in the wall), the profile will be a step function. As an initial approach this seemed promising, but later on we discovered that throughout the entire volume, too few fibres were detected with this approach. Instead we continued with 3D methods (See Sections 3.8.2 and 3.8.3).

4.5 Paper V: Ring-shaped object detector for non-isotropic 2D images using optimized distance transform weights

The method in Paper IV is only applicable to 2D images with square pixels, so here we extend it to rectangular shaped pixels. This is needed when the resolution is different in different directions, as in our case. You could avoid this by interpolating a set of virtual intermediate pixels, but this can be complex for the difficult shapes of the fibres and will also require larger images, thus longer processing times. As our main tool is based on the use of the DT, this can be avoided by optimizing the chamfer weights used in the computation of the DT for pixels of size $0.7 \times 5 \mu m^2$, instead of square pixels. After that the method described in Paper IV can directly be used. Comparing the results to an interpolated image, we note that although the resolution of the wall profile is much less in the original image, we reach quite satisfactory results even though the pixel size ratio is as high as seven. However, the same limitations as in Paper III still exist, i.e., the method only use 2D information. This is not a feasible approach when a large amount of fibres need to be detected.

4.6 Paper VI: Some measurements of fibres in volume images of paper using medial representations detected on the distance transform

Even if only some of the fibres are found by the method described in Section 3.8.3, these *are* individually segmented with a unique label and can be used when developing measurements for fibres in true 3D images derived from a real paper sample, this has never been done before. An important task is to select a suitable representation (data type) of the fibre. We use two, since it is advantageous to treat the fibre wall and fibre lumen separately. The wall is represented as a medial representation of centred voxels within the wall, producing an inner surface structure which is one voxel thick, following the original shape of the fibre wall. In each of these surface voxels, the distance to the closest void voxel is stored. This can be calculated by using DTs. The DT_{image} is iteratively thinned: one DT_{image} layer is removed at each iteration, while preserving the structure of the wall, i.e., without adding or removing cavities or tunnels in the surface. From this representation, we can then directly obtain estimates of, e.g., fibre wall thickness. The lumen representation is

created by a similar procedure, creating a 3D curve through the centre of the lumen with the radius to the closest wall voxel in each of the curve voxels. We can use the two representations to guide the computations of a number of measurements that will characterize each fibre: fibre wall volume, fibre lumen volume, fibre length, average fibre wall thickness, and degree of collapse. We continued developing more measurements in Papers VII and VIII. Maybe the most interesting aspect with these measures is the level of details possible to obtain. Even if we take into account the errors from the 3D reconstruction and segmentation, the amount of detail for each single fibre we manage to segment is extremely detailed.

4.7 Paper VII: Estimating fibre twist and aspect ratios in 3D voxel volumes

Continuing thinking about what to measure that is both possible to implement and useful for fibres in three spatial dimensions, we developed twist and aspect ratio measures. Both these measures definitely need to use information based on three spatial dimensions. Aspect ratio can be measured from cross-sections only, but will not give as accurate measures for an entire fibre, as a 3D analysis is capable of. Both measures rely on accurate cross-sections along the fibre length, thus we must estimate the local orientation, i.e. where the fibre is heading in the 3D voxel space, since we move along the fibre lumen curve. We fit a line to the local neighbourhood using least squares, which was robust enough for our case. Using this direction estimation as the normal of a local plane, we resample the voxels that intersect that plane to get a cross-section that is perpendicular to the local direction. Using principal component analysis of the wall voxels in this plane, we obtain a base vector along the axis of maximal variance and a third perpendicular vector. The local aspect ratio is now the width/height ratio in this local co-ordinate system. By averaging this ratio throughout the fibre, we get an accurate estimate of aspect ratio for a single fibre. The twist measure should not include the effect of fibre bending. To attenuate this effect, we project a vector of maximal variance from a neighbouring voxel onto a local plane that has its normal vector parallel to the estimated local orientation. To be able to assess at least a rough accuracy estimate, we constructed a set of simulated fibres with known geometrical properties and compared the twist and aspect ratios from the voxel volume to the exact one. The aspect ratio estimated was only a few percent from the exact measure derived from the spline parameters, which is good. Twist is a more complex property to measure. We believe that the 10–15% error our method has, is still within reasonable limits.

4.8 Paper VIII: Using distance transform based algorithms for extracting measures of the fibre network in volume images of paper

Our representations, introduced in Paper VI, were not fully exploited, so we continued with more measurements. Compared to Paper VI, we have added fibre slenderness, free fibre length, fibre curl, torsional resistance, and paper density. Slenderness measures the length to diameter ratio, curl how much the fibre swirls around (non-straightness), and torsional resistance the ability of the fibre to resist rotation during

twisting forces. Free fibre length is a network-level measure, thus relating properties from several neighbouring fibres. It measures the total length of the part of a fibre that does not touch another fibre. It can be seen as a density measure based on a fibres 3D position. Paper density was estimated by the “rolling ball” technique. Since the paper surface is quite complex, we need to smooth it to be able to set a fixed surface that differentiate between paper and non-paper. Rolling ball is implemented efficiently by a 3D DT, where the ball radius determines the level of smoothing. Essentially all measures can be divided into three classes, depending on their type: (1) paper level; (2) fibre level; and (3) network level. As noted, we have only one paper level measure (density), caused by the relatively small paper sample. Larger samples are needed to measure paper level properties. Fibre level is anything which characterizes single fibres. Since the fibres have quite complex shapes, there are many things to measure using 3D analysis. The network level, measuring how fibres interact has largely been left out because of the problem of finding enough fibres to sample the true network in a reasonable way.

4.9 Paper IX: 3D Pore distribution estimation in digitized paper samples

Paper is a porous media, where large parts of the interior are pockets of void. Most of our efforts have been focused on the fibres, but the void is also a network of complex shape that very much influences paper properties. In fact, it is a kind of “inverse fibre network”. An example of a 3D measure for the voids, is to determine their size distribution, i.e. a histogram of the void volumes, and if there are any systematic variation along paper thickness. To be able to measure individual void volumes, we need to segment the void and also define how to separate void regions that are interconnected but are reasonable to define as separate. We tried to use the 3D watershed transform for this purpose. To eliminate the large number of the small pockets, we pre-filter the image to close these. This is in fact a kind of weak smoothing. Moreover, we compared their approach to a 3D scanning method developed at PFI (the Norwegian Pulp and Paper Research Institute), which avoids the segmentation step and still can provide the average size distribution of pores. We were not able to accurately segment the pore network into individual regions. Most certainly an improved pre-processing for the watershed based segmentation, alternatively another more suitable approach will be needed. Inspired by the PFI approach of void analysis without segmentating individual pores, we introduce the idea of using the histogram of the DTs to compare the pore size properties of a paper sample. As no spatial localization is possible with this approach, we instead add a method to measure how the pore density (porosity) varies as a function of surface distance.

5 Conclusions and further results

The procedure we developed to digitize a sample of paper works well. It combines the unique properties of Scanning Electron Microscopes with the high-precision microtome cutter, to sequentially produce a stack of image slices from a paper sample, which we can reassemble into an accurate 3D digital voxel volume. The geometrical distortions are small, and furthermore, it is possible to digitize larger pieces of paper using the same approach. In this section, we will conclude with some reflections and future ideas.

The amount of accuracy required for registration is very high. This is caused by the fact that we are trying to simulate a 3D imaging device by capturing a sequence of 2D images and patch these together. Any misalignment during this process will undoubtedly introduce artefacts. They are visible as a kind of aliasing, seen as jerky/choppy translations of fibres as well as scattered material inside the paper when you browse through the volume looking at 2D cross-sections. Furthermore, the difference in resolution aggravates this problem, since we only have 1/7 of the xy -resolution along the z -axis. This is the cause of the large difference between adjacent slices. To be able to interpolate, we need to find the fibre cross-sections in order to distinguish which regions belongs to each other. Simple grey-level interpolation does not give satisfactory results in this case, since the differences in position are well above one voxel diameter for adjacent slices. If we only were supposed to measure fibres on one 2D image at a time, the geometrical errors would only influence the accuracy of our measurements, not the ability to find the fibres. To be able to have a good chance of finding fibres in 3D images, an accurate registration and mosaicing is crucial.

It is reasonable to not expect both $1\ \mu m$ resolution and to cover an entire A4 sheet of size $297 \times 210 \times 0.10\ mm^3$, when digitizing a paper. Just imagine capturing this sample at $1 \times 1 \times 1\ \mu m^3$ voxels: we would have $297,000 \times 210,000 \times 100 \approx 6$ Tbyte of data! And, as you may have realized by now, we would not only have the sheer load of data, but moreover to extract a good percentage of the fibres in a paper sample. This is a complex task that needs further study before it is reasonable to try on even larger samples.

The most difficult part of the project has been the segmentation of individual fibres. Perhaps the most annoying issue with 3D image analysis is the fact that you know you have an extreme amount of data, and you just want to abstract this information into a higher level, e.g., to find some fibres. By only browsing through the data and looking at the image slices, it is easy to get an impression that this should not be difficult, but we assure you, do not try this at home.

By adapting and combining two 3D image analysis methods developed for arteries segmentation (Ordered Region Growing and SeparaSeed) we developed a fibre tracking algorithm that can provide a good base for individual fibre segmentation, where there are high demands on accuracy and robustness is required. The experiments we did on 2D based methods shows that they are not as robust as our 3D based algorithm. By looking at the changes needed for the 3D based method, we want to point out that even if the medical imaging and paper analysis domains should and

can share methods successfully, it is not a straight-forward task to convert methods between these different areas. However, if we had found the ORG method earlier, the focus would definitely have changed to true 3D segmentation much sooner, since this method suits our problem well. The SeparaSeed has been developed at CBA, so it was easy to get good local support when we found good use of this nice algorithm in our project as well.

To be able to measure the fibre network accurately, we do need to improve the segmentation method, since a high percentage of the total amount of fibres is needed. In principle, these measures are not more difficult to implement than the measures based on single fibres. As can be seen in the published papers, a large amount of fully 3D measurements of fibres have been developed, and as far as we know, true 3D measures with this level of detail for fibres has not been done before.

Validation of achieved measurements is also an issue. Ideally, an alternative technique would be useful to have something to compare with. We have tried to the best of our knowledge to implement sound measures, but a thorough estimation of the accuracy they can provide would be needed before they are practically useful. Since our methods are focused on individual segmentation of approximately cylinder-shaped fibres, any paper grade that does not have this kind of fibres will not be possible to analyze. Comparing the light microscopy and the SEM datasets, the LWC paper presents a considerably harder individual fibre segmentation task than the milk-carton in the SEM volume. Even assuming we could register and assemble these into an accurate 3D voxel image, our individual fibre analysis would not be able to find as many fibres as we can in the SEM dataset. The type of paper grade selected do influence how well our methods perform. An improvement would be to incorporate additional segmentation methods to handle the other types of objects that do occur in digitized paper samples. This would allow a more complete analysis, although how far we can get in this digital conversion, and analysis process, is difficult to speculate about.

It should be obvious that the task of digitizing a piece of paper, involves quite a large effort. Although some prefer to denote the result of this digitization a model, we believe that identifying many of the individual fibres and classifying the remaining parts of the paper are important steps forward. Until this is successfully implemented we cannot really say it is a digital model of paper. Although many measures do not require individual labelled fibres, we hope that the work done in this Thesis show that by doing so, many more opportunities for measures will be available. The scale of this problem is easy to underestimate, since there are both many practical details during the programming that is essential, as well as the development of the image analysis methods.

Our current use of the SeparaSeed method assumes that the diameter of all the wall holes and “glitches” due to anisotropic sampling are less than the diameter of the nearby lumen region, otherwise the lumen region will “leak” into the paper void region. By tuning the way we define connectivity, e.g., combining our DT based measure with local measures of the fibre shape around the ORG path, it should be possible to steer the SeparaSeed to favour filling along the ORG path and not as fast in the perpendicular direction. The surrounding DT_{image} shell approach could be

improved to handle even more neighbouring fibre material, so that tightly clustered fibres may have a more accurate segmentation of the fibre wall.

Segmentation for fibres without lumen but with adjacent fibres and filler material has not been considered yet, since that is probably an even harder problem. As the ORG method seems robust, an idea would be to seed the start and end of a lumen-less fibre, either by hand or if possible by automatic seeding, then let the ORG trace out a path between these points. It would certainly be necessary to add different restrictions, as curvature limiting the path, or maybe fit a 3D spline, to guide the ORG method. It may prove possible, but as we essentially can have a large blob of fibre material at any place inside the paper, it may only be reasonable for fibres with relatively small contact area to adjacent material. The fibre material not successfully segmented could maybe be classified as *BLOB* and still contribute to the calculations. A more parallel approach of finding the fibres seems a good idea, as they can cooperatively help to resolve which part should belong to which fibre, a task that is less robust when only analyzing one fibre at a time, as we did. We estimate a subset of the total fibre network with our approach and it may be possible to iteratively find better solutions. After the first iteration we have a set of fibres with an estimate on their path through a 3D space and which voxels they are composed of. The next iteration could use this, and may improve the fibre segmentation.

Even if the development of methods to perform the network-level measurements is not more difficult, they require an improved segmentation that can detect a large percentage of the fibres, to be of any use. Developing a more automatic seeding procedure would be an improvement for the ORG based segmentation, since manually marking start and end points for all fibres would be a too tedious task for large data sets. Since everything computer related is getting faster, it seems reasonable to anticipate much larger data sets and then more automation is needed. Some sort of skeletonization of the non-fibre-wall material could provide a start set of voxels to use for seeding, if an intelligent selection could be made, which is certainly not a straight forward task.

It was more fun to do the segmentation based on a 3D voxel space for our fibre material, compared to analysis of 2D images, which was developing into more of a test of endurance than improving the code. Focusing on methods for three spatial dimensions seems the obvious way to go, since only then the advantage of a captured 3D data set is fully utilized.

The experiment we did on 2D Gradient Vector Flow (GVF) snakes was interesting and instructive. We feel that snakes do have a lot of potential for segmentation, if only it were possible to increase their robustness, maybe by a human-in-the-loop approach as done in the “Live Wire” concept (Falcão et al., 1998). It is an appealing concept, restricting local shape deformations, but still enables the surface of the snake to adapt to a given shape. Using surface snakes in the spatial dimensions to find the boundary between the lumen and the inner fibre wall surface, would be an interesting experiment. Still, our focus on discrete voxel based methods feels competitive, as they often are much more efficient than snakes.

There is ongoing work with simulations of mechanical properties of fibre networks, which could be combined with our approach of finding fibres in a real paper

sample, to enable advanced simulations of some of the mechanical properties of the paper. In the thesis of Heyden (2000), she uses a Finite Element Method (FEM) to estimate material properties of fibre fluff. This material is less dense than paper, but also has a complex network of fibres. FEM is a general method to simulate the behaviour of large models with properties governed by differential equations. Since the resulting equation system can be very large, especially for complex 3D models, a fast parallel computer may be required to get reasonable computation times. FEM simulations has been used with great success to estimate temperature distribution, fluid flow, mechanical deformations, and more. For instance, SAAB uses FEM based computer simulations for their 3D car models, to predict aerodynamic behaviour and also for collision simulations. By using a single 3D voxel fibre as geometrical input to a FEM solver, it would be possible to predict its behaviour during different kinds of load conditions. Simulating a large piece of the fibre network would give even more valuable information on the mechanical properties of the digitized paper sample.

Since the entire process of creating paper already is model-based, as each computerized process control equipment is based on its own model, constantly fine-tuned to perform better, it seems reasonable to continue and also model the end-product, the paper. Model based approaches are nice in the way they can bring order out of chaos. The simplifications introduced make the problem more manageable and still some of the main properties can be preserved. Our attempts can hopefully provide a starting ground for others interested in paper analysis, by pointing at the specific problems that need to be solved and why they are difficult.

The idea of converting a paper into digital format, really unlocks many possibilities of advanced paper analysis. This is of course biased by our preference for computer science and image analysis, but anyway ...



CALVIN AND HOBBS © Watterson. Reprinted with permission of UNIVERSAL PRESS SYNDICATE. All rights reserved.

References

- Adams, Rolf and Bischof, Leanne. Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 16(6):641–647, June 1994.
- Advanced Visualization Laboratory (AVL) – VolumePro graphics card, 2002. Available [Online]: <http://www.avl.iu.edu/technology/volumepro/> [19 October 2002].
- Bardage, Stig. Three-Dimensional Modeling and Visualization of Whole Norway Spruce Latewood Tracheids. *Wood and Fiber Science*, 33(4):627–638, 2001.
- Bardage, Stig and Daniel, Geoffrey. Three-Dimensional Analysis of the Collapse Behavior of Kraft-Cooked Norway Spruce Fibers. *Wood and Fiber Science*, 34(3):382–390, 2002.
- Bergholm, Fredrik. VISIT – VISual Information Technology, 2002. Available [Online]: <http://www.cb.uu.se/visit/> [19 October 2002].
- Beucher, Serge. Watersheds of functions and picture segmentation. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (CASSP'82)*, pages 1928–1931, Paris, 1982.
- Björck, Åke. *Numerical Methods for Least Squares Problems*. SIAM, 1996.
- Borgefors, Gunilla. Distance transformation in digital images. *Computer Vision, Graphics, and Image Processing*, 34:344–371, 1986.
- Borgefors, Gunilla. Hierarchical chamfer matching: A parametric edge matching algorithm. *Pattern Analysis and Machine Intelligence (PAMI)*, 10:849–865, November 1988.
- Borgefors, Gunilla. Applications using distance transforms. In *Proc. Aspects of Visual Form*, pages 83–108, Singapore, 1994.
- Borgefors, Gunilla. On Digital Distance Transforms in Three Dimensions. *Computer Vision and Image Understanding (CVIU)*, 64(3):368–376, November 1996.
- Borgefors, Gunilla and Svensson, Stina. Fuzzy border distance transforms and their use in 2d skeletonization. In *Proceedings 16th International Conference on Pattern Recognition (ICPR)*, volume 1, pages 180–183. IEEE Computer Society, August 2002.
- Brändström, Jonas. *Morphology of Norway Spruce Tracheids with Emphasis on Cell Wall Organisation*. PhD thesis, Swedish University of Agricultural Sciences, Uppsala, Sweden, February 2002.
- Canny, John. A computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, November 1986.
- Carneiro, Bernardo, Silva, Cláudio, and Kaufman, Arie. Tetra-Cubes: An algorithm to generate 3D isosurfaces based upon tetrahedra. In *Anais do IX SIBGRAPI*, pages 205–210, Brazil, 1996.

- Computer Graphics – many centres working with computer graphics can be found here, 2002. Available [Online]: <http://mambo.ucsc.edu/psl/cg.html> [19 October 2002].
- Danielsson, Per-Erik. Euclidian distance mapping. *Computer Graphics and Image Processing (CGIP)*, 14:227–248, 1980.
- European Synchrotron Radiation Facility (ESRF), 2002. Available [Online]: <http://www.esrf.fr> [19 October 2002].
- Falcão, Alexandre, Udupa, Jayaram, Samarasekera, Supun, Sharma, Shoba, Hirsch, Bruce, and Lotufo, Roberto. User-Steered Image Segmentation Paradigms: Live Wire and Live Lane. *Graphical Models and Image Processing (GMIP)*, 60(4):233–260, July 1998.
- Fayyazi, Arash. Texture based duplex-board layer segmentation, May 2002. Lic.Thesis, Linköping University, Campus Norrköping, ITN – Department of Science and Technology.
- Frank, Petra. A method to characterize the material distribution in paper. Technical Report UPTEC F00058, Uppsala University School of Engineering, Uppsala, Sweden, July 2000.
- Grevera, George, Udupa, Jayaram, and Odhner, Dewey. An Order of Magnitude Faster Isosurface Rendering in Software on a PC than Using Dedicated, General Purpose Rendering Hardware. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 6(4):335–345, 2000.
- Hamarneh, Ghassan. *Towards Intelligent Deformable Models for Medical Image Analysis*. PhD thesis, Chalmers University of Technology, Göteborg, Sweden, September 2001.
- Hansson, Peter and Johansson, Per-Åke. A new method for the simultaneous measurement of surface topography and ink distribution on prints. *Nordic Pulp and Paper Research Journal*, 14(4):315–319, 1999.
- Herman, Gabor. *Image Reconstruction from Projection: the fundamentals of computerized tomography*. Academic Press, New York, 1980. Ch. 10.
- Heyden, Susanne. *Network modelling for the evaluation of mechanical properties of cellulose fibre fluff*. PhD thesis, Lund University, Department of Mechanics and Materials, Sweden, 2000.
- Kaestner, Anders. *Non-Invasive Multidimensional Imaging Applied on Biological Substances*. PhD thesis, Halmstad University, Sweden, October 2002.
- Karlsson, Håkan and Fransson, Per-Ivar. STFI FiberMaster gives the papermaker new muscles, 2000. Available [Online]: <http://www.stfi.se/documents/contract/pulp/fibermart.htm> [19 October 2002].
- Kim, Bo Hyoung, Seo, Jinwook, and Shin, Yeong Gil. Binary volume rendering using the slice-based binary shell. *Manuscript available from the authors*, page 29 pages, 1999.

- Lacroute, Philippe and Levoy, Marc. Fast Volume Rendering Using a Shear-Warp Factorization of the Viewing Transformation. In *Proc. SIGGRAPH '94*, pp. 451–458, Orlando, Florida, July 1994.
- Levoy, Marc. The Digital Michelangelo Project, 2002. Available [Online]: <http://graphics.stanford.edu/projects/mich/> [19 October 2002].
- Lorensen, William and Cline, Harvey. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, July 1987.
- Lundqvist, Sven-Olof. Karakterisering av ved- och fiberegenskaper, 1999. PhD course at SFTI, in Swedish.
- Mähler, Anders. EuroFEX – the STFI research papermachine, 2000. Available [Online]: <http://www.stfi.se/documents/contract/papermaking/fexmachine.htm> [19 October 2002].
- Moëll, Mattias. *Digital Image Analysis for Wood Fiber Images*. PhD thesis, Swedish University of Agricultural Sciences, Uppsala, Sweden, December 2001. Silvestria 225.
- Mueller, Klaus and Yagel, Roni. Fast Perspective Volume Rendering with Splatting by Utilizing a Ray-Driven Approach. In *Proc. of the Visualization '96 conference*, pages 65–72, 1996.
- Niskanen, Kaarlo, Nilsen, N., Hellén, E., and Alava, Mikko. KCL-PAKKA: Simulation of the 3D structure of paper. In *The Fundamentals of Papermaking Materials, The Eleventh Fundamental Research Symposium*, pages 1177–1213, Cambridge, UK, September 1997. Pira International.
- NLM (National Library of Medicine) – the Visible Human Project[®], 2002. Available [Online]: http://www.nlm.nih.gov/research/visible/visible_human.html [19 October 2002].
- Nordin, Bo. *IPAD, version 2.0 & IMP — and IPAD application*. Centre for Image Analysis, Uppsala, Sweden, 1997. Internal Report No:6.
- Ourselin, Sébastien, Roche, Alexis, Subsol, Gérard, Pennec, Xavier, and Ayache, Nicholas. Reconstructing a 3D structure from serial histological sections. *Image and Vision Computing*, 19:25–31, 2000.
- POV-Ray[™] – The Persistence of Vision Ray Tracer, 2002. Available [Online]: <http://www.povray.org/> [19 October 2002].
- Pozo, Roldan. Template Numerical Toolkit (TNT), 2002. Available [Online]: <http://math.nist.gov/tnt/> [19 October 2002].
- Rosenfeld, Azriel. Digital topology. *American Mathematics Monthly*, 86(8):621–630, 1979.

- Rusinkiewicz, Szymon and Levoy, Marc. QSplat: A Multiresolution Point Rendering System for Large Meshes. In Akeley, Kurt, editor, *SIGGRAPH 2000, Computer Graphics Proceedings*, pages 343–352. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000. Download software [Online]: <http://graphics.stanford.edu/software/qsplat/> [19 October 2002].
- Samuelsen, Emil, Helle, Torbjørn, Houen, Per, Gregersen, Øyvind, and Raven, Carsten. Three Dimensional Imaging of Paper by Use of Synchrotron X-Ray Microtomography. In *Proc. TAPPI International Paper Physics Conference*, pages 307–312, San Diego, USA, 1999.
- Sonka, Milan, Hlavac, Vaclav, and Boyle, Roger. *Image Processing, Analysis, and Machine Vision*, chapter 3.3. PWS Publishing, 2nd edition, 1999.
- The Swedish Pulp and Paper Research Institute (STFI), 2002. Available [Online]: <http://www.stfi.se> [19 October 2002].
- Thibault, Xavier, Bloch, Jean-Francis, and Boller, Elodie. Felt structure characterised by synchrotron microtomography. *Appita Journal*, 55(2):145–148, March 2002.
- Tizon, Xavier and Smedby, Örjan. Segmentation with gray-scale connectedness can separate arteries and veins in MRA. *Journal of Magnetic Resonance Imaging*, 15(4):438–445, April 2002.
- Udupa, Jayaram and Herman, Gabor. *3D Imaging in Medicine*. CRC Press, 2nd edition, 1999.
- Udupa, Jayaram and Odhner, Dewey. Shell Rendering. *IEEE Computer Graphics and Applications*, 13(6):58–67, 1993.
- Udupa, Jayaram and Raya, Sai. Shape-Based Interpolation of Multidimensional Objects. *IEEE Trans. on Medical Imaging*, 9(1):32–42, 1990.
- Udupa, Jayaram and Samarasekera, Supun. Fuzzy Connectedness and Object Definition: Theory, Algorithms, and Applications in Image Segmentation. *Graphical Models and Image Processing*, 58(3):246–261, 1996.
- Vincent, Luc and Soille, Pierre. Watersheds in Digital Spaces: An Efficient Algorithm Based on Immersion Simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 1991.
- Wang, Huaijun and Shaler, Stephen. Computer-Simulated Three-Dimensional Microstructure of Wood Fibre Composite Materials. *Journal of Pulp and Paper Science*, 24(10):314–319, Oct 1998.
- Wilhelms, Jane and Gelder, Allen Van. Octrees for faster isosurface generation. *Computer Graphics*, 24(5):57–62, November 1990.
- Wyvill, Brian, McPheeters, Craig, and Wyvill, Geoff. Data Structure for Soft Objects. *The Visual Computer*, 2(4):227–234, 1986.

- Xu, Chenyang and Prince, Jerry. Gradient Vector Flow: A New External Force for Snakes. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 66–71, Los Alamitos, California, USA, June 1997.
- Yang, C.-F., Eusufzari, A.R.K., Sankar, R., Mark, Richard, and Perkins Jr., Richard. Measurements of geometrical parameters of fiber networks – Part 1. Bonded surfaces, aspect ratios, fiber moments of inertia, bonding state probabilities. *The Swedish paper magazine (Svensk papperstidning)*, 1978.
- Yim, Peter, Choyke, Peter, and Summers, Ronald. Gray-scale Skeletonization of Small Vessels in Magnetic Resonance Angiography. *IEEE Transactions on Medical Imaging*, 19(6):568–576, June 2000.

6 Other publications and conferences

The author has also contributed to the following publications:

- Fayyazi, A. and Aronsson M., “Towards Assembling a Small Digital Volume of Paper”, in Proc. “Microscopy as a Tool in Pulp and Paper Research and Development”, pp. 124–133, 1999, STFI, Stockholm.
- Aronsson M. and Fayyazi, A., “3D Fibre analysis of paper - a survey”, Centre for Image Analysis, Blue report No:30, ISSN 1100-6641, Uppsala 1999.
- Aronsson M. and Larsson K., “Titta inuti papper”, (In Swedish, English translation: “Look inside paper”), in “Nordisk Papper och Massa”, pp. 44–45, No:2, 2001.
- Aronsson M. and Svensson S., “Curvature measurements for fibres in 3D images of paper”, pp. 165–168, SSAB Symposium 2002, Lund.

The following conferences have been visited when the articles in this thesis were presented:

- “Microscopy as a Tool in Pulp and Paper Research and Development”, June 21–22, 1999, STFI, Stockholm, Sweden.
- “Swedish Society of Automated Image Analysis (SSAB ’00)”, Mars 7–8, 2000, Halmstad Univ., Sweden.
- “12th Scandinavian Conference on Image Analysis (SCIA ’01)”, June 11–14, 2001, Bergen, Norway.
- “Computer Vision and Pattern Recognition (CVPR ’01)”, Dec 9–14, 2001, Kauai Islands, Hawaii, USA.
- “Swedish Society of Automated Image Analysis (SSAB ’02)”, Mar 7–8, 2002, Lund, Sweden.
- “International Conference on Pattern Recognition (ICPR ’02)”, Aug 11–15, 2002, Québec, Canada.
- “International Conference on Image Processing (ICIP ’02)”, Sep 22–25, 2002, Rochester, New York, USA.

7 Acknowledgements

Gunilla, my supervisor, who has supported me during these five years of image analysis studies. At times when the results of some experiments turned out to be zilch, she was quite calm and did not advertice for a new PhD student. Thanks for the trust. Ewert, if you ask me, VISIT was a great idea. But you also managed to convince others before I started as a PhD student in the VISIT programme, so thank you, and SSF, for letting me focus on research and not fundings. Lennart and Fredrik, former and present VISIT program directors. Örjan and Olle H., for creating *the* image data set, and also helping me out with the pulp and paper problems. Stina, if it were not for you, I probably had not finished by now, since I would still be debugging my gfx-code instead of focusing on fibres :-). Thanks a lot for the co-operation, the push forward, and also for redefining my scale about efficient people. Arash, for our joint work in the “3D tracking of fibres” project. Björn, that you together with Gunilla and Örjan initiated the 3D fibre project, which Arash and I would become a part of. The reviewers of this thesis. Thanks to you (Catherine, Geoffrey, Gunilla*, Hans, Ingela⁺, Mats, Stina⁺, Örjan). You helped me correct the many less optimal :-) formulations and descriptions in the text, and without your comments this would have been a much less readable text. Note the “plus i kanten”, which is only earned by outstanding thesis-support achievements, by returning nearly as many pages of possible suggestions of improvements as original thesis pages I sent to them! Gunilla always scrutinizes all of the papers intended for publishing, thus often improving them considerably. Much appreciated. Rune & Co at PFI, for sharing their data and some nice meetings. Joanna, for creating a very difficult data set for us, thus making us realize that we need to learn more. Ida, for our common work on paper analysis. Olle E., for ensuring that my computer behaves, and also for not bugging me why I need more than 1 Gbyte of swap space, but instead upgrading the hardware. Good computer support is essential, and you certainly provide that. Jocke, for keeping me awake at our monday seminars by constantly asking interesting questions to the speaker. Bosse, not for the IMP environment, sorry, I do like MATLAB better, but for fun discussions about computer related stuff. Ola, for being supportive and encouraging, but also for letting me test drive his new car so that I got another reason to continue in the CBA-gambling-team. Mats, for inadvertently reminding me that I should (start to) exercise a bit more, as his way-over-normal *km* biking count is very impressive. My biking between ICA Kniven and home need extended floating point arithmetics to describe the mileage-ratio to Mats. My room mate, Felix, for nice company in room 123 and elsewhere, but also for keeping the plants alive. Lucia, both for improving my english and making me realize that it is quite normal to fly all over the globe just to present a poster. Lena, for making exceptionally good cakes at the monday seminars, but also by eliminating all kinds of practical problems. All other people at CBA, which makes up a very friendly and cosy environment. My best friends, Roger and Mikael, for beeing generally great and fun to be with. My parents, Margot and Harald, and my two sisters Cecilia and Ann-Sofi, for always supporting me in all possible and impossible ways. Bill Watterson, for your great Calvin and Hobbes comics, that can improve my mood in a split second. Jan Berglin, your comics about perfectly normal human behaviour is just awesome! SAAB, for making nice cars. Everyone else that I should have remembered to thank,

but did not. Thank you! So this is it, the final sentence of my thesis, I will not say a single word more about it, just that ...

HEY! DON'T UNPLUG MY COMPUTER.....



CALVIN AND HOBBS © Watterson. Reprinted with permission of UNIVERSAL PRESS SYNDICATE. All rights reserved.