

Segmentation of Wood Fibres in 3D CT Images Using Graph Cuts

Erik L.G. Wernersson, Anders Brun, and Cris L. Luengo Hendriks

Centre for Image Analysis, Swedish University of Agricultural Sciences
Box 337, SE-751 05 Uppsala, Sweden
{erikw, anders, cris}@cb.uu.se
<http://www.cb.uu.se>

Abstract. To completely segment all individual wood fibres in volume images of fibrous materials presents a challenging problem but is important in understanding the micro mechanical properties of composite materials. This paper presents a filter that identifies and closes pores in wood fibre walls, simplifying the shape of the fibres. After this filter, a novel segmentation method based on graph cuts identifies individual fibres. The methods are validated on a realistic synthetic fibre data set and then applied on μ CT images of wood fibre composites.

Keywords: microtomography (μ CT), graph cuts, wood fibres, composite materials.

1 Introduction

This article addresses the problem of segmenting individual wood fibres in synchrotron radiation micro computed tomography (CT) images of wood fibre composites, which is solved in a two step procedure. The first step is to identify the void inside wood fibres called the lumen. It is a challenging task, particularly when the fibre wall does not separate the lumen from the background. This is the case when cracks have formed, or when the membranes of the fibre pores are damaged or thinner than the maximal resolution of the imaging modality, see Fig. 1. In the second step, the segmented lumens are used to guide the identification of individual fibres using a graph cut segmentation algorithm.

Wood fibres are relatively cheap to produce and thus interesting to use as bulk material in, for example, wood fibre reinforced composites. The wood fibre is a tree cell that mainly consists of cellulose (40 – 44%), hemicellulose (15 – 35%) and lignin (18 – 35%) dry weight [1]. In a tree, it functions as a transporter of nutrition and water and is also the building block that gives wood materials their mechanical properties. It has an elongated shape with small pores, as seen in Fig. 1. However, during the manufacturing of bulk material, the cell structure may be damaged. The fibres sometimes break into small pieces and cracks appear in the fibre walls. This gives the material new mechanical properties and it introduces significant difficulties for segmentation algorithms.




Fig. 1. (A) A CT image with a cross section of a fibre with a pore in the middle. (B–D): Three hollow wood fibres extracted from a CT volume with the methods described in this article. A pore can be seen in the wall of fibre (D).

3D images of fibrous materials can be acquired in high resolution using confocal microscopy of 2D slices cut from the sample with a microtome and by CT of full 3D samples [2,3]. With the microscopy approach the resolution is seldom isotropic while the CT technique yields isotropic images and has also the advantage of being non destructive. Acquisition time varies from about one week for microtomy down to 20 minutes with the CT technique.

Wood fibre composites typically have a tight fibre network with many points of contact, which makes identification of individual fibres challenging. A complete segmentation is however necessary to understand the mechanical properties of the composites. One characteristic of the fibre that usually can be distinguished or identified is the lumen. We therefore follow the approach of first identifying individual lumen, and from there identify the corresponding fibres.

1.1 Previous and Related Work

Existing lumen segmentation methods are based on geometrical approaches and are designed for binarized 2D slices where each pixel is either fibre or background. In [4], lumens are identified as non fibre objects that have a certain area, ellipticity and mean curvature. This method is later used in [5] and [6], where a total number of ten fibres are segmented by applying this technique to sequential slices of a volume image. Pores in the cell walls do, however, result in many failed segmentations using this approach. To solve the problems with pores, Aronsson carried out some experiments with active contour models, snakes, steered by Gradient Vector Flow [7] but concluded that it was not very useful since the lumen often are far from elliptical in 2D slices. Hagen and Holen [8] developed a method that was later refined and automated by Bache-Wiig and Henden [9] to handle pores. Their method is based on a series of operations involving the constrained distance transform [10] and can not operate on single image slices. It is further restricted to small pores in fibres that shouldn't be much deformed.

A good lumen segmentation can be used as a seed for segmentation of individual fibres. Aronsson [4] has presented a method for this in which the constrained 3,4- chamfer distance transform is calculated from each lumen in the fibre material. The fibre wall thickness is then estimated from the histogram of the distance




Fig. 2. (A) A synthetic image from [4] and (B) the corresponding segmentation result. (C), the result of our graph cut segmentation method. (D) A close up of the three fibres to the bottom right from (left) our method (right) Aronsson’s method.

values. Based on that, a segmentation is performed. The method works under the assumption of uniform fibre wall thickness and was demonstrated on a synthetic image that we have reproduced in Fig. 2. Since the fibre wall thickness is not even, there is a need for other methods.

Another approach to fibre segmentation is by tracking, i.e. by non global methods that follow individual fibres from seed points. There is a method based on the Radon transform by Axelsson [11], a region growing approach by Svensson and Aronsson [6] and a maximally stable extremal regions based method by Donoser and Bischof [12]. We aim at global segmentation of full volume images, while the tracking methods aims at segmenting one fibre at a time and often require user interaction.

1.2 Contributions

A filter that identifies and closes pores in the fibre walls is presented, which has the advantage over previous methods that it works on a wide range of shapes and sizes of pores and cracks. A new method, based on graph cuts, to segment individual fibres in volume images is introduced. It has the advantage over the methods described above that the restriction to even fibre walls thickness is removed. Evaluations of methods designed to operate on images of wood fibres are hard to perform, since no ground truth usually exists but here we describe preliminary results on how to generate realistic synthetic 3D images of wood fibres for evaluation.

2 Methods

In this paper, we propose a filter that closes pores in cross sectional 2D images of fibrous materials, which is based on curvature calculations of the fibre wall

boundaries. It identifies the opposite sides of pores as maxima in the curvature and connects suitable pairs of such points by line segments. We then segment individual fibres from 3D volumes using the lumen segmentation and graph cuts. The method transforms the volume image to a weighted graph structure and solves the segmentation problem in that domain before reconstructing the segmented volume image.

2.1 Closing Pores

In cross-sectional images of wood fibres, it is observed that the points of highest curvature at the boundaries are those at opposite sides of pores. Identifying pairs of points with high curvature and connecting them by line segments will thus close the pores.

The input data are binarized slices $I(x, y) \in \{0, 1\}$ extracted from 3D volumes. For each object in the image, boundaries are extracted by using 8-connectivity (4-connectivity would give a more rugged boundary) and the pixel coordinates are stored in the arrays $\hat{x}(c)$ and $\hat{y}(c)$, where c is discrete. Anti clockwise direction is defined as positive on the outside of objects and the coordinate arrays are treated as cyclic. The coordinates are upsampled by a factor three and smoothed by a Gaussian kernel, G , with $\sigma = t/2$ to reduce unwanted details and digitization effects. The parameter t is the expected fibre wall thickness. This gives $b(c) = (G*\hat{x}(c), G*\hat{y}(c))$. We denote the first derivative $\hat{x}'(c)$, and the second derivative $\hat{x}''(c)$.

The curvature κ at a point c of a curve is a measurement on how bent the curve is [13]. As a comparison, a circle with radius r has a curvature of $1/r$ everywhere. In our case, the anticlockwise boundary extraction gives a positive curvature where the border is convex. Hence the curvature will have positive maxima at the opposite sides of the pores. We also use two distance measures for the identification of pores, the shortest distance along the boundary (the geodesic distance) $d_G(b(i), b(j))$ and the Euclidean distance $d_E(b(i), b(j))$.

However, to connect any two points where the curvature has a local maximum is not selective enough. First of all, if the fibre wall has a thickness, t , then the curvature at the pore is $2/t$ and we demand the curvature to be at least half of that. We do not let the distance along the boundary be longer than the maximal circumference of a lumen m_c and we do not let the Euclidean distance be larger than half of that. We also want the ratio between these distance values to be strictly different from 1 to avoid connecting corners of fibre walls with polygonal shape. We also not allow the straight lines between the points to overlap fibre material by more than four pixels.

Taking these considerations into account, we connect suitable points with line segments using Bresenham's algorithm [14]. We illustrate the method on a CT image of a fibre cross section in Fig. 3 and the complete algorithm is described in Algorithm 1. Note that the algorithm is performed in all three coordinate directions in the volume to achieve more directional invariance.

```

Data: (A 2D binary image)
Result: (A 2D binary image where pores are closed)
for each object in the image do
  extract the border to  $(\hat{x}, \hat{y})$  ;
   $\tilde{x} = \hat{x} * G, \tilde{y} = \hat{y} * G, b(c) := (\tilde{x}(c), \tilde{y}(c))$  ;
   $\kappa(c) = \frac{\tilde{x}'(c)\tilde{y}''(c) - \tilde{x}''(c)\tilde{y}'(c)}{(\tilde{x}'(c)^2 + \tilde{y}'(c)^2)^{3/2}}$  (calculate the curvature);
  for all  $c_i, c_j$  do
    if  $\kappa(c_i), \kappa(c_j) \geq \frac{1}{t}$  and  $\kappa(c_i), \kappa(c_j)$  are local maxima;
     $d_G(b(c_i), b(c_j)) < m_c$  and  $\frac{d_G(b(c_i), b(c_j))}{d_E(b(c_i), b(c_j))} \geq 1.2$  and  $\int_{b(c_i)}^{b(c_j)} I(x, y) ds < 5$ 
    pixels then
      | connect  $b(c_i)$  and  $b(c_j)$  using Bresenham's algorithm;
    end
  end
end

```

Algorithm 1. The algorithm that closes pores




Fig. 3. (A) Cross section of a fibre with a pore. (B) There are many maxima in the curvature but in this case only three are strong enough and thus potential end points for lines that could close the pore. The location of those points are marked with black dots. (C) The line connecting the maxima at the gap is automatically selected and thus the fibre wall can be completed and the lumen correctly identified as the black region. More examples can be seen in Fig. 7.

2.2 Segmenting Lumen

The fibre walls will separate lumens from background when the pores are closed and thus, lumen can be segmented by labeling all non-fibre objects in 2D images. The standard procedure is then to classify the regions as background or lumen according to a few geometric features [4]. In our experiments, we simply classified objects as lumen or background depending on their area.

2.3 Automatic 3D Segmentation of Individual Fibres

It is possible to distinguish between fibre material and background in 3D CT images and from now on we also assume that individual lumen are segmented. When wood fibres are randomly distributed in some material, they have larger contact areas to their own lumen than to other fibres. As a consequence, when

```

Data: vol (a 2D or 3D fibre image);
lum (a lumen segmentation of vol);
capMap (an edge capacity map, see the text);
Result: seg (A segmentation of the fibres in vol)
for each label in lum do
  | construct the graph, each pixel is a node and neighboring pixels share an
  | edge with capacity according to capMap.;
  | connect all voxels of the label to the source and remaining lumen voxels to
  | the sink. ;
  | solve the min-cut optimization problem and store in MC ;
  | for Each node in MC do
  | | if the node is connected to the sink then
  | | | set the corresponding pixel in seg to the current label;
  | | | remove the corresponding pixel in vol and lum;
  | | end
  | end
end
set all voxels in seg that corresponds to any lumen in lum to zero;

```

Algorithm 2. The fibre segmentation algorithm. It uses a lumen segmentation as seed and divides the volume with graph cuts.




Fig. 4. After the image (A) is binarized into (B), each of the 25 object pixels is replaced by a node connected to its neighbours as in graph (C). The orange pixels in (B) represent lumen and we connect the left one to a special node, called the source and the right one to the sink, *t*. The minimal cut in the graph (C) is marked with a dashed line and separates the two fibres.

modeling the flow or transfer of something, say water, between different lumen in a fibre network; the fibre to fibre contact will limit the flow. Hence, to find such regions and cut there would also be to segment the individual fibres. Such modeling exists in graph theory and the solutions are called minimal cuts [15].

A volume image can be converted to a graph by letting each voxel be a vertex and connect each adjacent voxel by an edge. When constructing a graph, each edge is allowed to have an individual (directed) capacity. We take advantage of that in two ways. First, we must assure that the flow from small lumen is not restricted because the number of pixels around it increases with the distance. This is solved by letting the capacities of the edges decrease linearly with the distance from the lumen until a distance corresponding to the largest expected

fibre width. We let the capacities increase linearly with the distance from the boundary between fibre and background. We call the result a capacity map for the graph structure. The full 3D algorithm is described in detail by Algorithm 2 and a 2D illustration is shown in Fig. 4.

3 Data

To validate our algorithms, they were tested on three types of data: synthetic 2-D images from a previous publication [4], realistic synthetic 3-D datasets of fiber materials recently developed in our lab and data from fibers scanned at ESRF in Grenoble, France.

3.1 Realistic Synthetic 3D Fibre Volumes

To validate our segmentation algorithm, we constructed a program to generate synthetic images of fibres in a volume with the dimensions $100 \times 100 \times 200$, where the fibres are aligned roughly along the last dimension. This is the case in paper and in many composite materials.

Each fibre is individually shaped from a set of variables that are uniformly distributed within an interval. The outer rim of the fibre cross section is modeled by a superellipse, $((x/r_1)^n + (y/r_2)^n)^{1/n} \leq 1$, where $n \in [2, 4]$. The two principal diameters of the ellipse are $r_1, r_2 \in [6, 16]$, the thickness of the fibre wall is $t \in [1.3, 2]$ and the fibre length is $L \in [100, 160]$. The bending of the fibre is modeled by aligning the fibre along a circle of radius $R \in [150, 300]$. The fibre is also twisted around the longitudinal axis, in total $f \in [0, 2\pi]$ radians, and randomly rotated around the same axis to make sure the bending of the fibre is applied in a random direction each time. Variations in the surface texture of the fibre were modeled by Perlin noise [16] and pores were added randomly to model irregularities and membranes in the fibre wall.

The fibres were stacked together by a simple procedure that placed each fibre inside the volume at a random position where it did not cover any other fibre, and then it was shifted by random steps in the two directions perpendicular to the fibre direction until it either reached the border of the volume or hit another fibre. In this manner, most fibres have contact with each other. Finally, the rendering of the fibre volume allowed us to extract different kinds of volume data: an anti-aliased gray-value model of the fibre volume similar to the CT images of fibres that we acquired; a thresholded image that has similar characteristics to the thresholded data sets we use for segmentation purposes; and finally, a volume where each voxel is labeled according to which fibre it belongs to. We also had the option to generate these images with or without pores, which was particularly useful to validate our pore-closing algorithm.

3.2 CT Images

The 3D CT images were imaged at the European Synchrotron Radiation Facility in Grenoble with a voxel size of $0.7 \times 0.7 \times 0.7 \mu\text{m}$ after the reconstruction. For




Fig. 5. A synthetic image is displayed in (A). A lumen extraction without closing the pores in the fibre walls gives the result shown in (B). If the proposed method is used, a better result, shown in (C), is obtained.



Fig. 6. Segmentation of fibres from the artificial volume with the proposed graph cut method. (A) The original volume from which lumen were segmented. (B): Each fibre is segmented, labeled and displayed.

our purposes, they are binarized by first applying a SUSAN filter [17] and then thresholding segmentation. Fibre material take up $\approx 17\%$ of the volumes. Parts of three extracted fibres can be seen in Fig. 1.

4 Results

The described methods were implemented for MATLAB (Release 2008b, The MathWorks) and MEX using C++ and run on a PC with 8GB RAM under Red Hat Linux. The running time for the pore closing in Fig. 5 was less than 10 minutes. Segmentation of the fibres in Fig. 2 took less than a minute and for the 3D problem in Fig. 6 less than 5 minutes.

4.1 Closing of Pores

We let the method process each slice of the synthetic 3D volume in Fig. 5 where the fibre walls are full of pores. There are in total 144141 voxels that correspond to lumen in that volume (the number approximated by analyzing the same volume before the pores were created). Results from the experiment that show that the method is successful are displayed in Tab. 1.



Fig. 7. Results of the fibre closing method on cross sections of fibres from CT-images

Table 1. Evaluation results of the pore closing method on a synthetic 3D image. Missing voxels are those that should have been classified as lumen but were not. Surplus voxels are pixels falsely classified as lumen. Max overlap is the number of pixels with fibre material that lines connecting pores are allowed to pass through.

	No Closing	With Closing				
Max overlap	-	2	3	4	5	inf
Missing voxels	20891	16375	9599	8382	8572	9377
Surplus voxels	0	119	415	753	934	1055
Total pixel errors	14%	11%	7%	6%	7%	7%

Results of the closing method on single fibre cross sections from μ CT data is shown in Fig. 7.

4.2 Individual Fibre Segmentation

We constructed the graph representations using 4- connectivity for plane image slices and 6- connectivity for volumes and used the MAXFLOW software [18] to solve the min-cut optimization problem.

The segmentation method was tried on the same 2D synthetic image as the method of Aronsson, see Fig. 2, where we first dilated each lumen as far as possible before it hit the background. In that case, our method is better since it does not leave any unassigned pixels and it resolves boundaries between fibres with less bleeding.

We have used the method on 3D images as well. For a synthetic volume, the segmentation worked excellently as shown in Fig. 6. We are also able to extract fibres from CT images, a few are shown in Fig. 1.

5 Discussion and Conclusion

The lumen closing method works best when the fibre direction is close to orthogonal to the slicing plane but is quite robust as seen in the example in Fig. 3 where the fibre clearly is not orthogonal to the plane. The method works on a wide range of shapes as seen in Fig. 7. The fibre closing can easily be extended to non-isotropic pixel sizes. The missing voxel errors in Tab. 1 are mainly due to multiple cracks in the fibres, in the same slice. The surplus voxels are created

where the fibres are densely packed and small areas between them are classified as lumen. From Tab. 1, we conclude that the method works as expected. We think that the remaining 6% errors are due to pores larger than the fibre wall or when there are many of them in a single cross section.

The fibre segmentation will return one individual fibre per supplied lumen. Hence, good lumen segmentation is important.

From experience, we know that fibres show a great variation in morphology and at this early stage of the development of a synthetic dataset generator we are only able to model some aspects that are relevant for this particular segmentation task. The building of a more complete environment for simulation of wood fibre materials is a task for future work within the scope of our current research project.

Acknowledgments

The authors would like to thank Karin Almgren, Maria Axelsson, Gunilla Borgefors, Kristofer Gamstedt, Bjørn Steinar Tanem, Catherine Östlund and all other persons in the WoodFibre3D project. The CT images were acquired at EFSR in Grenoble, France. This work was funded through WoodWisdom-Net under project number 352006A.

References

1. Bowyer, J.L., Shmulsky, R., Haygreen, J.G.: *Forest Products and Wood Science, an introduction*, 5th edn. Blackwell Publishing, Malden (2007)
2. Donoser, M., Wiltsche, M., Bischof, H.: A new automated microtomy concept for 3D paper structure analysis. In: *Proc. of the IAPR Conf. on Machine Vision Appl.*, pp. 76–79 (2005)
3. Samuelsen, E., Houen, P.J., Gregersen, O.W., Helle, T., Raven, C.: Three-dimensional imaging of paper by use of synchrotron x-ray microtomography. In: *Proc. Int. Paper Physics Conf.*, pp. 307–312 (1999)
4. Aronsson, M., Borgefors, G.: 2D segmentation and labelling of clustered ring shaped objects. In: *Scandinavian Conference on Image Analysis*, pp. 272–279 (2001)
5. Lunden, J.: *Image analysis methods for evaluation of fibre dimensions in paper cross-sections*. Master's thesis, Royal Institute of Technology, Stockholm, Sweden (2002)
6. Svensson, S., Aronsson, M.: Some measurements of fibres in volume images of paper using medial representations detected on distance transforms. In: *IEEE Computer Society Conf. Computer Vision Pattern Recognition.*, vol. 2 (2001)
7. Aronsson, M.: *On 3D Fibre Measurements of Digitized Paper*. PhD thesis, Swedish University of Agricultural Sciences (2002)
8. Holen, R., Hagen, M.: *Segmentation of absorption mode x-ray micro tomographic images of paper*. Master's thesis, Norwegian University of Science and Technology, Trondheim (2004)
9. Bache-Wiig, J., Henden, P.C.: *Individual fiber segmentation of three-dimensional microtomograms of paper and fiber-reinforced composite materials*. Master's thesis, Norwegian University of Science and Technology, Trondheim (2005)

10. Piper, J., Granum, E.: Computing distance transformations in convex and non-convex domains. *Pattern Recognition* 20(6), 599–615 (1987)
11. Axelsson, M.: 3D tracking of cellulose fibres in volume images. In: *IEEE Int. Conf. on Image Processing*, vol. 4, pp. 309–312 (2007)
12. Donoser, M., Bischof, H.: Efficient maximally stable extremal region (MSER) tracking. In: *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 553–560 (2006)
13. Carmo, M.P.D.: *Differential Geometry of Curves and Surfaces*. Prentice Hall, Englewood Cliffs (1976)
14. Bresenham, J.E.: Algorithm for computer control of a digital plotter. *IBM Systems Journal* 4(1), 25–30 (1965)
15. Sonka, M., Hlavac, V., Boyle, R.: Chapter 7.6. In: *Image Processing, Analysis, and Machine Vision*, 3rd edn. Thomson Engineering (2008)
16. Perlin, K.: An image synthesizer. *SIGGRAPH Comput. Graph.* 19(3), 287–296 (1985)
17. Smith, S.M., Brady, J.M.: Susan - a new approach to low level image processing. *Int. Journal of Computer Vision* 23(1), 45–78 (1997)
18. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE PAMI* (9), 1124–1137 (2004)