

Automatic Analysis of Log End Face Images in the Sawmill Industry

Kristin Norell

*Faculty of Forest Sciences
Centre for Image Analysis
Uppsala*

**Doctoral Thesis
Swedish University of Agricultural Sciences
Uppsala 2010**

Acta Universitatis Agriculturae Sueciae
2010:25

ISSN, 1652-6880
ISBN, 978-91-576-7502-6
© 2010 Kristin Norell, Uppsala
Tryck: Universitetstryckeriet, Uppsala University, 2010

Abstract

Kristin Norell. *Automatic Analysis of Log End Face Images in the Sawmill Industry*.
Doctoral Thesis.

ISSN, 1652-6880, ISBN, 978-91-576-7502-6.

At present grading of sawlogs in a sawmill relies on visual inspection wherein a human expert grades a log every few seconds as it passes on a conveyor belt. This tedious and difficult work is prone to substantial inter- and intra-grader variability.

This dissertation presents methods for automatic analysis of end faces from Scots pine (*Pinus sylvestris* L.) and Norway spruce (*Picea abies* (L.) Karst). Two features are extracted: the pith (centre core) position and the number of annual rings. The pith detection uses local orientations to estimate the centre of the annual rings in a manner that is robust to disturbances as knots and cracks as well as partial coverage of dirt or snow. The number of annual rings is counted using the polar distance transform, a tool developed here. This transform combines the image intensity and the circular shape of the rings so that the annual ring pattern can be outlined in rough and noisy images. First the marks on end faces from uneven sawing are removed using an automatic method developed in this work.

The data are images of untreated end faces mostly acquired at sawmills. A large amount of the data was imaged using a camera mounted above a conveyor belt at a sawmill, collecting images every month during one year. In total, the data consists of over 4000 images of pine and spruce. In this dissertation an algorithm for generating synthetic log end face images is also presented. The synthetic data can be used as a tool for developing image analysis methods.

The annual ring measurements were thoroughly evaluated on pine end face images acquired using the mounted end face camera. This evaluation shows that the method performs equally well as an experienced manual grader for grading the logs into quality classes. The method can thus be used as a component of an automatic grading system, overseen by a manual grader.

Key words: image analysis, polar distance transform, coniferous end faces, wood quality, pith detection, annual ring measurements, synthetic wood images

Author's address: Centre for Image Analysis, Box 337, SE - 751 05 Uppsala, Sweden.

List of Enclosed Papers

The dissertation is based on the following papers, which are referred to by Roman numerals in the text.

- I Kristin Norell, Creating synthetic log end face images, In proceedings of the 6th International Symposium on Image and Signal Processing and Analysis, Salzburg, Austria, pp. 353-358, IEEE, 2009
- II Kristin Norell, Joakim Lindblad, and Stina Svensson, Grey weighted polar distance transform for outlining circular and approximately circular objects, In proceedings of the 14th International Conference on Image Analysis and Processing, Modena, Italy, pp. 647-652, IEEE, 2007
- III Robin Strand and Kristin Norell, The polar distance transform by fast-marching, In proceedings of the 19th International Conference on Pattern Recognition, Tampa, FL, IEEE, 2008
- IV Kristin Norell and Joakim Lindblad, Spatially-variant morphological operations on binary images based on the polar distance transform, In proceedings of the 19th International Conference on Pattern Recognition, Tampa, FL, IEEE, 2008
- V Kristin Norell and Gunilla Borgefors, Estimation of pith position in untreated log ends in sawmill environments, *Computers and Electronics in Agriculture*, 63(2):155–167, 2008.
- VI Kristin Norell, An automatic method for counting annual rings in noisy sawmill images, In proceedings of the 15th International Conference on Image Analysis and Processing, Salerno, Italy, pp. 307-316, Lecture Notes in Computer Science 5716, Springer Verlag, 2009
- VII Kristin Norell, Counting annual rings on *Pinus sylvestris* end faces in sawmill industry, Submitted for journal publication

For all papers, except Paper III, the work was mainly performed by the author. For Paper III the contribution was mainly on the application and experiments as well as parts of the writing.

Related work

In addition to the papers included in the dissertation, the author has also contributed to the following publications:

Kristin Norell, Finding the Pith in Rough Cross-Sections of Logs in Sawmill Environments, Swedish Symposium on Image Analysis, pp. 129-132, Umeå, 2006

Kristin Norell, Joakim Lindblad and Stina Svensson, The polar distance transform, Swedish Symposium on Image Analysis, pp. 67-70, Lund, 2008

Contents

Automatic Analysis of Log End Face Images in the Sawmill Industry

1	Introduction	9
1.1	Objectives and motivation	9
1.2	Organization of the dissertation	10
Part I: Background		11
2	Sawlog quality	13
2.1	Tree growth	13
2.2	Wood degradation	15
2.3	End face analysis	16
2.4	Quality measurements at Swedish sawmills	18
3	Real data acquisition	21
3.1	Image acquisition	21
3.2	Real data	27
4	Synthetic data	39
4.1	Background	39
4.2	Generating 2D end face images	40
4.3	Additional experiments - 3D sawlogs	45
Part II: Theory		47
5	Image analysis	49
5.1	Pixelwise operations	49
5.2	Operations on pixel neighbourhoods	52
6	Distance transforms	59
6.1	Computations using chamfering	59
6.2	Computations using fast marching	61
6.3	The polar distance transform	64
6.4	Applications using the polar distance transform	68
Part III: Applications		71
7	Preprocessing	73
7.1	Determining pixel size	73
7.2	Automatic cropping	79
7.3	Automatic removal of marks from uneven sawing	85

8	Analysis of end face features	89
8.1	Pith detection	89
8.2	Annual ring detection	102
9	Concluding remarks	107
9.1	Conclusions	107
9.2	Future work	108
	Summary in Swedish	111
	References	113
	Acknowledgements	117
	Errata to Paper III	119

Papers I–VII

1 Introduction

This dissertation is the result of five years of work at the Centre for Image Analysis in Uppsala, Sweden. The project was performed in collaboration with SDC - IT company for the Swedish forestry sector, with the goal to develop methods useful in the sawmill industry for analysing digital images of coniferous log end faces.

1.1 Objectives and motivation

Wood features are measured both in the forest industry and in the field of dendrochronology. Image analysis can be a helpful tool in such measurements, since manual inspection is often tedious and difficult, and requires experts to perform it. With automatic image analysis methods, the results can be more consistent than for manual measurements, especially if these are performed by different persons.

The applications of wood measurements determine possibilities and limits for end face preparation and image acquisition. In this work, the application is quality analysis at sawmills, where features of the end face are used already today as a parameter for grading each log into a quality class. At present, the grading at Swedish sawmills is performed manually by a grader who has many different features to take into account. If parts of this analysis can be done automatically, possibly with support from the grader, the demands on the grader are less, and the consistency of the result will probably be improved. Image acquisition and analysis must be carried out during approximately two seconds, if the quality measurement is not to become a bottle-neck at the sawmill. Here, images acquired with digital cameras with light in the visible wavelengths, were used. This gives a fast acquisition and also rather inexpensive hardware with little effect on the sawmill environment. The data has been collected at sawmills, using both system cameras and a camera more suitable for industry applications, that was mounted above a conveyor belt where logs were transported and imaged.

Imaging end faces at sawmills means that there are usually no possibilities for treating the end faces in any way before acquisition. The untreated end faces require methods for analysis that are robust to the rough surface that is the result after harvesting, storage and transport of the logs from the forest to the sawmill. End faces are occasionally also partly covered with dirt or snow. Methods previously developed for analysis of higher quality end faces, such as in dendrochronology, where end faces are often treated (sanded and possibly oiled) and imaged through a microscope or scanner could therefore not be applied with good results on the images studied here.

In this dissertation, two different features have been measured on coniferous log end faces: the pith position and the number of annual rings. In the development of the annual ring measurement method a new theoretical tool, the polar distance transform, was introduced. This transform combines the circular shape of the rings with the intensity values of the image, so that it is possible to outline the rings. A method for automatic removal of marks on the end faces that occur due to uneven sawing has been developed in this work and was applied to the images before the

feature extraction. An algorithm for generating synthetic end face images was also developed. The synthetic data resembles pine end faces and the simulation includes wood features and properties occurring from sawing, storage, camera setup and image acquisition.

The methods for pith detection and annual ring measurements were developed and evaluated on real data acquired at sawmills. The work was performed on Scots pine (*Pinus sylvestris* L.) and Norway spruce (*Picea abies* (L.) Karst), which will here be referred to as pine and spruce, respectively.

1.2 Organization of the dissertation

The work has been divided into three different parts in this dissertation: background, theory and applications, respectively. The background includes a description of sawlog quality, such as the procedure of tree growth, wood degradation, and quality measurements at Swedish sawmills. Imaging techniques are also described together with the cameras that were used during the work and the acquired data. The algorithm for creating synthetic data is also included in this part (Paper I).

The second part describes theory, including basic image analysis concepts that were used in the developed methods or in other similar applications. Furthermore, a chapter on distance transforms is included, in which the polar distance transform is described (Paper II, Paper III, and Paper IV).

The third part of the dissertation is devoted to applied end face analysis. The first chapter presents preprocessing of the data such as computing pixel size in the images and a method for automatic cropping of a large dataset of images. These parts have not been published earlier, and are included in the dissertation for completeness. The preprocessing also includes the description of a method for removing marks across the end face that occurred due to uneven sawing (Paper VII). In this part of the dissertation, the analysis of end face features are also described: detection of pith position (Paper V) and computing the number of annual rings (Paper VI and Paper VII). New results are presented for both methods, using data that was not analyzed earlier. Conclusions and a summary in Swedish close the dissertation.

Part I

Background

2 Sawlog quality

The quality of a sawlog relies to a high extent on the quality of the tree from which the log originates. Many sawlog features are related to tree growth such as straightness, number of knots or branches, amount of compression wood, and annual ring width, while other features rather relates to wood degradation, e.g., presence of rot and blue stain. Also harvesting and handling can cause damage to the log, which leads to the occurrence of scars or cracks. In this chapter the growth and degradation of a conifer is described on the macro level with focus on features that are visible on an end face. Sawlog quality measurement at Swedish sawmills is also briefly described.

2.1 Tree growth

As a tree grows from a small plant to a large tree the stem grows higher and wider, the bark thickens and new branches are formed. In the innermost part of the stem is the centre of the annual rings, the *pith*.

2.1.1 Annual rings

The stem thickens as new cells are created between the outermost, already existing wood cells and the bark. Due to different growth conditions such as temperature, light, and nutrition, cells will have different properties. In conifers, good growth conditions result in cells with thin cell walls and large voids (lumens). During bad conditions cells will instead grow thick walls and small lumens (Bowyer et al., 2007). The better and worse conditions coincide with the early and late part of the growth season, respectively, and the different types of wood created are thus referred to as earlywood and latewood.

If a tree stem is cut and the cross section is examined with the naked eye, earlywood appears brighter than latewood. The cells created during one growth season are thus represented by a bright region of earlywood and a dark region of latewood. Hence each set of earlywood and latewood together represents an *annual ring*, or growth ring. The latewood region is often thinner than the corresponding earlywood region since growth is slower during late season.

Examples of earlywood and latewood cells in pine are shown in Figure 1. The lumens are seen as black areas, cell walls are grey and the light grey areas are lignin, which is present both in the outer parts of the cells and in-between cells. The cells have been growing from left to right in the image, starting with earlywood cells with large lumens and thin cell walls, growing towards latewood. The transition from earlywood to latewood is smooth and growth is slower and slower during the late season. Finally, to the right in the image, one growth season has passed, and a new period of good growth conditions starts, with earlywood cells produced. The transition between latewood and earlywood is rapid. The image is captured by a camera attached to an epifluorescence microscope.

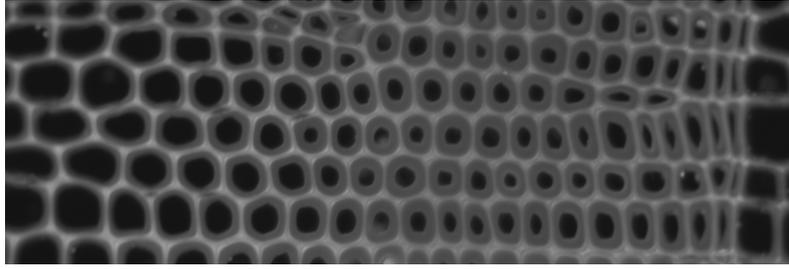


Figure 1: Wood cells from pine seen on micro level with earlywood to the left, growing through a smooth transition towards latewood. To the right, the sharp transition from latewood to earlywood is visible. The piece of wood is approximately $26.2 \times 8.7 \mu\text{m}$. [Image acquisition: Bettina Selig]

2.1.2 Heartwood

When the tree grows older, properties of the innermost wood cells in the stem will change. These cells get a high concentration of certain chemicals that give them different properties compared to regular cells. A region with such cells is called *heartwood* and appears in the centre part of the stem, as opposed to the *sapwood* surrounding it. The border between heartwood and sapwood is not generally restricted to a specific annual ring. Heartwood presence is not completely determined by tree age, but has for pine been shown by Kärenlampi and Riekkinen (2002) to be related to tree maturity, which is a measure inversely proportional to area increase of stem cross section over time. Pine heartwood can often be seen as a region with a darker or more reddish colour than surrounding wood. This is illustrated in Figure 2 where the red area in the centre is heartwood. In spruce, there is no distinct colour difference between heartwood and sapwood and therefore the heartwood is more difficult to detect visually.

2.1.3 Branches and knots

The branches are created during tree growth starting from the pith. Conifer branches are formed in whorls every year in the top of the stem. As the tree continues to grow, the branches are embedded in the stem by new wood cells. This embedding can create perturbations of the annual ring pattern in the region close to the knot. In Figure 2, a knot is visible at the boundary between heartwood and sapwood. A close look at the area around the knot shows the perturbed annual ring pattern. Since branches are formed in whorls it is not unusual that an end face have several visible knots.

2.1.4 Compression wood

Occasionally trees are influenced by different types of mechanical stress, e.g., hard wind constantly coming from one direction, or a growth position where the tree is leaning towards light. If this is the case, reaction wood is formed to compensate for the mechanical stress. In conifers the reaction wood is present on the lower side of



Figure 2: A pine end face showing heartwood, a knot, and blue stain.

a leaning stem or a branch and called *compression wood*. These wood cells have very thick cell walls and tensions which is a problem in wood usage. When lumber or planks are dried after sawing present compression wood will most likely give bent products, since different parts will shrink to different amount and the tensions will cause deforming. Due to the thick cell walls, compression wood is darker, more reddish than regular wood and the latewood area is often wider than usual. Compression wood on an end face is often limited to a rather small region of a few annual rings, but can in severe cases cover larger areas.

2.2 Wood degradation

Wood degradation can occur both during growth and after harvesting. There are different kinds of degradation, such as rot, staining and insect attacks. Here, the focus is on rot and staining which are both caused by fungi. Fungi spread by spores through air, water, or insects. As the spores come in contact with wood mycelium starts to grow and eventually a fruiting body might form. All wood fungi feed from wood cells, but in slightly different ways. How the fungi affect wood quality thus depends on what is left in the cell after the attack.

2.2.1 Rot

Fungi causing rot are of a few different kinds, e.g., *brown rot* and *white rot*, feeding on different components of the wood cells. Brown rot feeds mainly on cellulose, leaving the lignin in the cells unaffected. Lignin is brown, and thus the result after this attack is a brown region. White rot, on the other hand feeds on both cellulose and lignin, which means that the brown lignin is removed, thus the wood is bleached. Both of these types of rot cause severe strength reduction of the wood.

2.2.2 Staining

Other types of fungi mainly stain the wood, without causing large effects on wood strength. These fungi feed on material from cell lumens, not affecting strength as much as rot fungi. Therefore, the decay is not as severe as in the cases with rot and the effect is more of an esthetical matter. The staining is often dark blue, thus called *blue stain*. Fungi are transported to the wood through the air and they attack the end face of a log or the parts of a stem or log that are not covered by bark. After the attack mycelium grows through the sapwood, but leaving the heartwood unaffected. The pine end face shown in Figure 2 is attacked by blue stain, which is seen as a blue discolouration of the sapwood.

2.3 End face analysis

In this work, the expression *end face* refers to a cross-section surface that is visible at the end of a log or a wood disc. End faces can be of different quality with varying readability, in the range from sanded and oiled end faces via end faces cut with a sharp tool in laboratory, freshly cut trees in the forest up to logs harvested, stored in the forest, transported on a truck and then stored yet some time at a sawmill. The difference between a freshly cut end face and one stored at a sawmill is shown in Figure 3, where the same end face is imaged twice: first after sawmill storage and then after sawing off a thin disc with the only purpose to produce a newly sawn end face. The difference between the two end faces is mainly the roughness of the surface and the marks created from uneven sawing during harvesting in the upper image. Note that the end face is rotated approximately 180 degrees between the two images. The end face is imaged using the same camera in similar conditions only a short time apart.

The expected readability of an end face thus differs a lot, depending on how the end face was treated. Heartwood can generally be difficult to detect since it does not always have a discriminating colour compared to the sapwood. For a well prepared end face, sanded and possibly oiled, most of the mentioned features can be seen with the naked eye, i.e., pith position, annual rings, knots, compression wood, rot, and blue stain. However, if annual rings are very narrow or affected by rot they can be difficult to distinguish. For end faces in the other end of the scale, as they appear in sawmill industry, it can be harder to detect all mentioned features. The ring pattern is not always clear and a rough surface can hide slight cases of rot and compression wood. In very rare cases, often with severe rot, even the pith can be difficult to detect.



Figure 3: A pine end face imaged as found at the sawmill (upper) and shortly afterwards when newly sawn (lower).

Table 1: Rules for grading spruce sawlogs. Details regarding the different concepts can be found in Anon. *VMR* (2007).

	Grade	
	1	2
Knots, whole mantle surface	Max 60 mm regardless of knot type.	Spike knots max 120 mm. Other knots unlimited.
Growth rings	Min 12	-
Straightness	Max 20 cm loss of saw yield.	Max 120 cm loss of saw yield.
Indication of top rupture	Not allowed	Allowed
Blue stain	Not allowed	Allowed
Open scar	Scar which affects the scaling cylinder is not allowed	Depth of scar max 20% of scaling cylinder diameter.
Bark-enclosed scar	Length max 2× top end diameter	Allowed
Forest rot	Not allowed	Max 5% of log end surface

If an end face is imaged, the imaging itself influences readability. Imaging concerns all techniques from capturing images through a microscope, desktop scanner, a CT scanner to a digital camera, which all give different types of images and varying image quality. Illumination and image resolution are also factors to consider.

The difference in end face quality and imaging techniques requires image analysis of end faces to be tuned for the specific application. The features analyzed may be similar, but methods suitable for one application, does not always work well for another. In this dissertation end faces from sawmill environments are analyzed, i.e., end faces from logs that were harvested in the forest, stored for some time, transported to the sawmill and then stored again. Imaging has been performed with digital cameras of different kinds, using light in the visual spectrum.

2.4 Quality measurements at Swedish sawmills

SDC - IT company for the Swedish forestry sector, determines standards and guidelines for measuring sawlog quality in Sweden. Different features are used to grade each log into a quality class, where spruce and pine sawlogs are separated into two and four different classes, respectively. The rules for grading are listed in Table 1 and Table 2 (Anon. *VMR*, 2007). Table 1 was modified after personal communication with Jacob Edlund, SDC. The highlighted rows describe the rules regarding annual rings, which is the grading feature that has been considered in this dissertation.

The measurement area for annual rings is a specific part of the end face: the interval 20 to 80 mm from the pith in the direction that gives the lowest result, i.e.,

Table 2: Rules for grading pine sawlogs. Details regarding the different concepts can be found in Anon. *VMR* (2007).

Log type	Grade			
	1	2	3	4
Log type	Butt log	Not butt log	All log types	All log types
Knots, whole mantle surface	Max 20 mm, all knot types. Max 5 knots.	Sound knots max 120 mm Other knots max 60 mm		Spike knots max 120 mm. Other knots unlimited.
Knots within 15 dm from butt end	-	A minimum of two distinct whorls or one sound knot.	-	-
Knot swelling	Max 5	-	-	-
Growth rings	Min 20	-	Min 12	-
Straightness	Max 20 cm loss of saw yield.			Max 120 cm loss of saw yield.
Indication of top rupture	Not allowed			Allowed
Blue stain	Not allowed			Allowed
Forest rot	Not allowed			Max 5% of end surface

where the mean annual ring width is the largest, but not affected by knots. Measurements are performed at the low end of the log, except for spruce butt logs (i.e., logs from the lowest part of the stem), which are measured in top end. Choosing the correct measurement area is important, both when it comes to distance and direction from the pith. Letting the end face in Figure 2 again illustrate a concept shows that the rather eccentric pith means that the average annual ring width differs above and below the growth centre. It is thus clear that to determine the correct number of annual rings requires the correct measurement area to be chosen.

For spruce sawlogs twelve rings are required in the measurement area for the highest quality, i.e., grade 1, while the number of rings is irrelevant for grade 2, as described by Table 1. For pine sawlogs (Table 2) two different thresholds affect the grade. For the highest quality, 20 rings are required, while it for grade 3 is enough with twelve. For grades 2 and 4 the number of annual rings is irrelevant. The pine end face shown in Figure 3 has 44 rings in the measurement area and thus belongs to grade 1, if taking into account only the number of annual rings.

When referring to the quality grades further on in the dissertation, only the number of rings will be considered. Because of this, and to make the grades correspond

to the same number of rings for spruce and pine, another quality grading will be used. Here, grade *C* corresponds to a number of rings below twelve; *B* is twelve or above for spruce, and from twelve to 19 for pine; *A* is above twenty, and will be measured only for pine. With this notation, *C* is the worst quality class for both species, and the highest measured class for spruce is *B*.

At present, the end face analysis at Swedish sawmills is performed manually by a log scaler or grader. Typically, the grader is positioned next to a conveyor belt where the logs are passing, either lengthwise or sidewise. The log scaler grades each log according to the rules, where one of the grader's tasks is to take into account the number of annual rings in the measurement area. Often a screen with a video or an image of the current end face is a help in this analysis.

3 Real data acquisition

The data used in this dissertation are images of log end faces essentially acquired at sawmills. This chapter describes the end faces, the different cameras that were used for acquisition and the resulting images. First, the procedure of image acquisition and the concept of a digital image are briefly described. Digital images in two dimensions are handled throughout the work hence such images are the focus here.

3.1 Image acquisition

To capture a digital image is to project a piece of the three dimensional world onto a two dimensional picture, sampled in space, intensity, and time. Here, images with sample points in a square grid are considered. The intensity is sampled in each point, corresponding to the incoming light at a specific time.

An image sensor captures light reaching it through the camera. The amount of light is determined by the lens aperture and the exposure time, where the first regulates how much light passes through the lens. A large aperture gives a bright image, but it also means a short depth of field, which is the distance from the camera where the imaged objects appear sharp. The exposure time determines for how long the image sensor integrates light, thus the higher the exposure time the brighter is the image. However, a long exposure time can give motion blur in the image, if the object is moving during acquisition.

3.1.1 *Imaging sensors*

A digital imaging sensor consists of several sensor elements, each collecting a signal representing the incoming light in a small area. The light is transformed into an electrical charge and converted to a voltage. Two common sensor types in digital imaging are the charge-coupled device (CCD) and the complimentary metal-oxide semiconductor (CMOS) sensors, which uses different techniques.

The CCD sensor is a shift register which moves the captured signals across the sensor before storage. The measured electrical charge in a sensor element is not converted into a voltage until it reaches one corner of the sensor where it can be read out. This makes it impossible to read out the information from all sensor elements simultaneously, thus readout must be performed for one row at a time. An advantage with this solution is that no extra space is needed on the sensor for conversion of the signals, thus the space can be used for integrating light, i.e., the fill factor is high. A schematic illustration of the CCD-sensor is shown in Figure 4, where information from the third row is being read out. Each charge is transported on the sensor towards the bottom row, and then out to the right.

In the CMOS sensor the conversion from electric charge to voltage is performed in the sensor element itself using transistors. Making room for the conversion on the sensor often results in lower fill factor compared to the CCD, since the light collecting element must be smaller. The fact that each element has its own conversion means that the readout can be performed in random order, or simultaneously. On

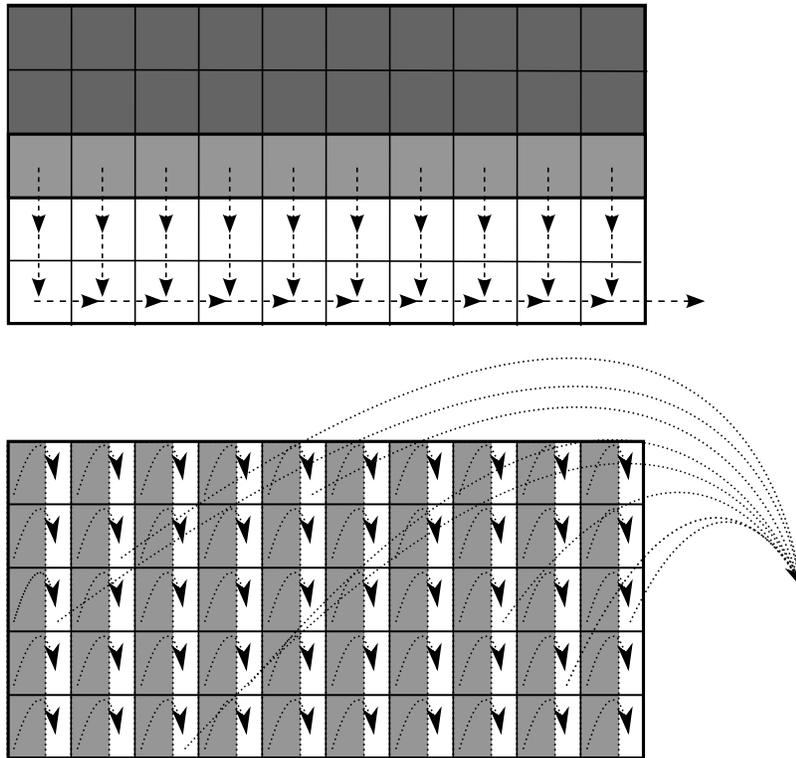


Figure 4: Schematic illustration of the readout from a CCD sensor (upper) and a CMOS sensor (lower). Dark areas represent elements that were already read out, grey areas are elements currently read, and white illustrate areas used in the read-out. The arrows illustrate how the information is transported.

the other hand, simultaneous readout might not be possible due to the large amount of information that must be handled if the number of sensor elements is high. A schematic illustration is shown in Figure 4. Here, the information is first processed at each sensor element and then read out. In the illustration all pixels are active simultaneously, and the readout is shown for a few random pixels.

3.1.2 Image size and resolution

When visualising a captured intensity image each sample point is represented by a picture element, *pixel* covering a small area of the image. The total number of pixels determines the *image size*, such as 3872×2592 pixels or 10.2 mega pixels. Image size can be equal to the number of sensor elements in the camera, but it can also be modified by resampling. This is the case, e.g., when a preview of a captured digital image is shown on the display of a compact camera. The preview is a smaller, sub-sampled, version of the captured image. The word image also denotes a sub-image, i.e., a part that has been cut out from a larger image.

Image size must not be confused with *image resolution*, which defines how close two point objects can be positioned and still be resolved as two objects in the image. The resolution depends on several things e.g., sampling frequency and how the incoming light is spread in the camera optics. According to Nyquist's sampling theorem the sampling frequency must be at least the double of the spatial frequency of the signal that is sampled to capture all details (Unser, 2000). High sampling frequency is, however, not enough to accomplish high resolution. The optics must also be of high quality so that the incoming light is not blurred in the camera lens. If the high frequencies are removed already in the camera optics, a high sampling frequency does not improve the image resolution.

Image resolution and sampling frequency are closely related to *pixel size*, a measure of how large area one pixel covers on the imaged object. If the objects in an image are positioned at different distances from the camera, the pixel size can vary in the image. It can also vary for a single object due to e.g., perspective distortion. The pixel size is often given in length units, e.g., mm or μm , representing the length of one side of a pixel.

3.1.3 Intensity and colour

In intensity images, or grey scale images, each pixel has a value proportional to the signal captured by the corresponding imaging sensor element. By mapping the pixel values on a grey scale from black to white the different intensities can be interpreted visually. In this work images with intensities between zero and 255 have been used, i.e., values that can be represented by 8 bits in a computer.

The wavelengths of the light determine the information that is acquired. Light in the visual spectrum give images that are similar to our way of seeing the surroundings. By varying the illumination and, if required, also the image sensors, the information captured from an object can vary. One example is infrared (IR) light, which was used in Gjerdrum and Høibø (2004) to determine the amount of heartwood on pine end faces. The different water content in heartwood and sapwood made them distinguishable.

One method to capture information from different wavelengths is to filter the light before it hits the sensor. The captured signal will then correspond to a specific wavelength band. To capture several images using different filters is called multi-spectral imaging. There, specific materials can be identified by their specific spectra, i.e., the captured intensities in different wavelengths. This was used by Nyström and Hagman (1999) to detect compression wood in planks of spruce. As a part of the analysis two classes were created representing earlywood and latewood, respectively. However, no information about the possibility to discriminate between these, and thus separate earlywood from latewood, was presented.

Intensities from different wavelengths can be combined to represent colour. This can be achieved with the mentioned filtering, but more common is to let different sensor elements capture light from specific wavelengths and create a colour image from these. This is the technique for the *Bayer filter* (Bayer, 1976) used in many cameras. Wavelengths corresponding to red, green, and blue light are captured by

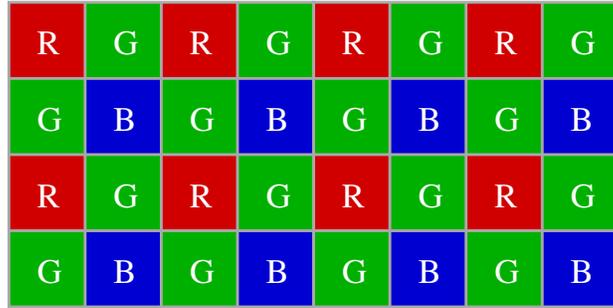


Figure 5: Illustration of a Bayer filter pattern where each sensor element captures light corresponding to wavelengths for red, green or blue light (R, G, and B).



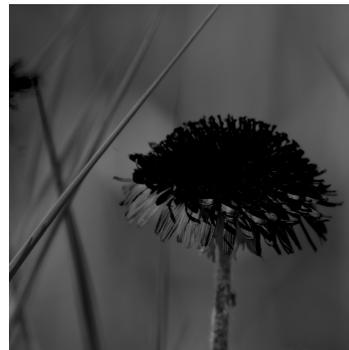
(a) Colour image



(b) Red channel



(c) Green channel



(d) Blue channel

Figure 6: A colour image and the three colour channels.

different sensor elements and then combined to produce a colour image. The pattern of filters used is shown in Figure 5 for a small part of an image. The fact that green is the dominating colour is an adaptation to the human visual system, with high sensitivity to the corresponding wavelengths.

After image acquisition, each pixel has a value corresponding to the incoming light in one of the three wavelength bands, or colour channels. To create the full colour image the missing information in an element is computed by interpolation from the neighbouring pixels. After interpolation, or *demosaicking*, colour is represented in each element by the three channels.

The difference between capturing three different images with red, green and blue filters, respectively, and to capture one using Bayer filtering is that with the latter technique, each colour is sampled at lower spatial frequency, and thus the resolution is lower. The missing information is interpolated to other elements, but the effect is that each colour channel is slightly smoothed. At positions with rapid colour changes the interpolation can therefore create artifacts that would not occur otherwise. This is referred to as *colour noise*.

Figure 6 shows a colour image captured using Bayer filtering. This image is built up from the three colour channels, red, green and blue and is referred to as an RGB-image. The RGB-space is a common output from a camera with each colour channel being an intensity image by itself. The colour channels are shown in Figure 6 with dark pixels corresponding to low numeric values and bright pixels corresponding to high values, respectively. It is clear that the yellow dandelion has high intensity both in the red and green channels, while it is close to zero in the blue one. For the grey background, the colour is more similar in all considered wavelengths.

There exist several other colour spaces where it is possible to represent a colour image, e.g., the *HSV-space*, separating between hue, saturation and intensity, or the *YIQ-space*. The different channels of the HSV and YIQ-space for the dandelion image are shown in Figure 7. The conversions from RGB to these spaces will not be described here, but the figure illustrates that the content in each image is highly dependent on the chosen conversion. The saturation component of the HSV-space is used in Cerda et al. (2007) for detection of annual rings. However, in the images analyzed in this work, the saturation was not discriminating the earlywood and latewood well enough. Here, the V-channel of the HSV-space was used in Paper V and a modified version of the Y-channel in the YIQ-space in Paper VI and Paper VII. The conversion from RGB-space to these colour bands are described in Section 5.1.



(a) Hue



(b) Y-channel



(c) Saturation



(d) I-channel



(e) Value



(f) Q-channel

Figure 7: The colour channels in the HSV-space (left column) and the YIQ-space (right column) for the image shown in Figure 6.

3.2 Real data

The images used in this dissertation have been captured using three different cameras. The imaging sensor types differ, as does the size of sensors and images. This information is listed in Table 3 for the cameras used and further on in the dissertation they will be referred to as Nikon, Canon, and PixeLINK, respectively. Nikon and Canon are consumer cameras while PixeLINK is designed for industrial applications. All cameras use a Bayer filter to capture colour images as described in Section 3.1.3. In images captured using the Nikon or Canon cameras, a small reference was imaged together with the end faces for computation of pixel size in the images. This reference can be seen in the upper part of Figure 10 located centrally. In principle, the reference can also be used for colour standardization to compensate for different illuminations and camera settings. This was not, however, necessary in this work.

Table 3: Camera specifications

Camera	Sensor type	Sensor size [mm]	Image size [pixels]
Nikon Coolpix 880	CCD	7.2×5.3	2048×1536
Canon EOS 400D Digital	CMOS	22.2×14.8	3888×2592
PixeLINK PL-A782	CMOS	7.7×10.5	2208×3000

The real data have been collected at different sites and times and with different aims. The overall focus has been to capture images of end faces as they appear at sawmills. This ensures that imaged end faces have realistic features and difficulties, which is important in the development and evaluation of image analysis methods. All datasets except for one dataset have been collected at sawmills. All are collected in Sweden and each dataset is described in detail below. If no imaging method is described, images were captured outdoors as the logs were stacked, with the camera mounted on a tripod. An overview of the datasets is found in Table 4.

3.2.1 Forssjö set

A small dataset of pine was captured in June 2005 at BooForssjö sawmill in Forssjö. End faces belonging to 12 logs were chosen in order to represent several different features. All logs had been rejected at the sawmill due to low quality and had therefore been stored for months outdoors. The low quality and long storage time together with choosing end faces with varying features resulted in a set of end faces where all are affected by rot, compression wood, or blue stain.

A disc approximately 30 to 40 mm thick was sawn from each log and the newly sawn side was sanded, after which both sides were imaged. The newly sawn side has been used only as reference. The discs were imaged on a table with the Nikon camera mounted at $h = 680$ mm above it, resulting in a pixel size that is approximately 0.3 mm. Two 25 kHz fluorescent lamps were mounted next to the camera, approximately $d = 200$ mm apart. Each end face was imaged twice using two different light temperatures: 3000 K and 4000 K. The image content was essentially the same for the images captured with different colour temperature, therefore only

Table 4: Overview of the collected datasets.

Dataset	Camera	Species	Number of end faces	Approximate pixel size [mm]
Forssjö	Nikon	pine	12	0.3
Ultuna	Nikon	spruce/pine	12/8	0.2
Nyby2005	Nikon	pine	25	0.2 to 0.8
Heby2006	Nikon	spruce	24	0.3
Heby2007	Canon	spruce	199	0.2
NybyOnline	PixeLINK	pine	> 4000	0.4

images with temperature 4000 K were used in further analysis. A schematic view of the setup is shown in Figure 8.

A typical end face from the Forssjö set is shown in Figure 10. An area affected by rot is visible to the upper right of the pith, while the dark region in the centre of the end face is heartwood. This disc contains compression wood which is visible on the newly sawn side, but not on the side shown in Figure 10.

3.2.2 *Ultuna set*

This dataset includes end faces that were imaged repeatedly each week during July to November 2005 at a test lot in Uppsala. End faces are from ten logs of about one meter length that were collected at a felling area and transported to the lot. The dataset includes six spruce logs and four pine logs that were all imaged from both sides, i.e., in total 20 end faces. At the time for log collection four spruce end faces were affected by rot, and two spruce and four pine end faces were affected by blue stain. The rest, i.e., six spruce end faces and four pine end faces showed clearwood. The logs were imaged using the Nikon camera and the pixel size is approximately 0.2 mm. A sample spruce end face with clear wood imaged in September is shown in Figure 11. The end face shown in Figure 2 is also a sample end face from the Ultuna set.

3.2.3 *Nyby2005 set*

The Nyby2005 dataset was captured at Setra Group Nyby sawmill in Björklinge in October 2005 using the Nikon camera. It consists of 25 pine end faces selected from logs both rejected and accepted at the sawmill measurement station. The distance between camera and object varied during image acquisition, resulting in pixel size between 0.2 and 0.8 mm. The end faces were chosen to include a variety of features such as rot, knots, eccentric pith, compression wood, dirt, double pith, and scars from felling. An image from the Nyby set is shown in Figure 12. The end face is severely affected by rot, which is visible as a rough area covering approximately two thirds of the end face.

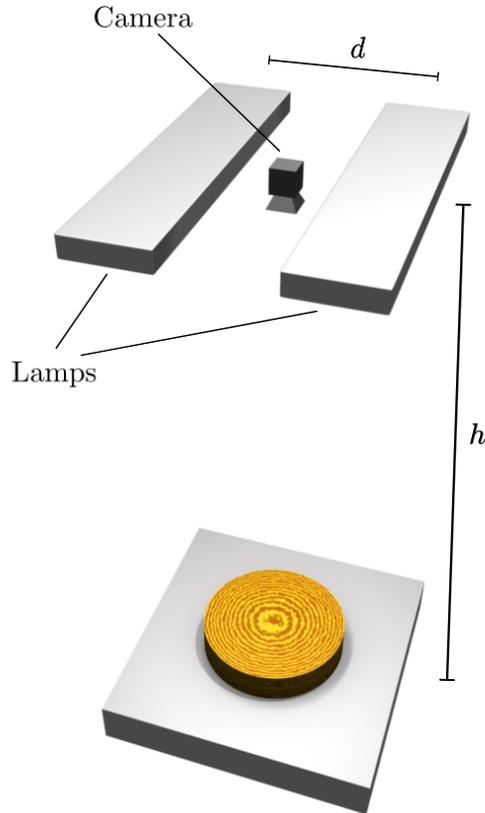


Figure 8: Camera setup for the Forssjö dataset. [Illustration: Patrik Malm]

3.2.4 Heby2006 set

With the Heby2006 set, the aim was to collect images of logs with sparse rings. Images of 24 spruce end faces were captured at Setra Group Heby sawmill in Heby in April 2006 using the Nikon camera. Pixel size was approximately 0.3 mm.

The end faces were classified by a grading inspector, thus divided into the different quality classes regarding the annual rings as described in Section 2.4. The dataset includes 20 end faces with grade *B* and four end faces with grade *C*. An end face with eleven rings in the measurement area, thus with grade *C*, is shown in Figure 13. The measurement area is to the upper right of the pith.

3.2.5 Heby2007 set

This dataset was also captured at Setra Group Heby sawmill in Heby, but in March 2007. Imaging was performed using the Canon camera and the pixel size is approximately 0.2 mm. This time, an effort was made to pick random end faces from

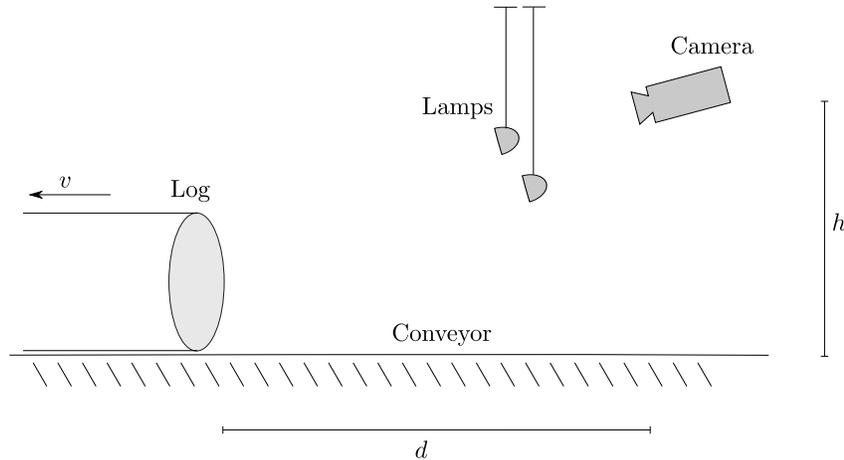


Figure 9: Camera setup for the NybyOnline dataset

different piles of logs to get a representative distribution of different features. In total 199 end faces were imaged and grading was performed by a grading inspector in the same way as in the Heby2006 set. Only six end faces have grade C . All other have twelve or more annual rings and were thus graded B . An end face with grade B is shown in Figure 14.

3.2.6 NybyOnline set

The NybyOnline dataset is the largest one. Images were captured monthly in on-line sawmill production at Setra Group Nyby sawmill in Björklinge during the year 2008 and beginning of 2009. The images were captured with the PixeLINK camera mounted at the sawmill measurement station. With this setup end faces were imaged as they passed on a conveyor belt, as illustrated in Figure 9. The camera was placed above the conveyor, depicting the logs as they moved away from the camera. It was positioned at the height $h = 0.9$ m above the conveyor and the distance $d = 2.2$ m from the end face in imaging position. The camera captures a region of approximately 0.7×0.6 m in 1800×1536 pixels, thus giving a pixel size of approximately 0.4 mm. A photoelectric cell was used to determine when to capture the image.

The PixeLINK camera has two alternatives for pixel readout: *rolling shutter* or *fast reset shutter*. With the rolling shutter a few pixel rows are exposed and read out simultaneously, meanwhile leaving the rest of the pixels inactive. With this technique there will be a small time difference between capturing the different pixel rows. If the object moves during acquisition it results in motion blur. With fast reset shutter, all pixels are activated, or reset, simultaneously, but the readout will still be performed for a few rows at a time due to the large amount of data that needs to be stored. The rows that have not yet been read out are still active, integrating light. This technique is intended to be used together with a flash or an external shutter that can steer the amount of light falling into the sensor.

In the NybyOnline setup no external shutter was used, and it was impossible to use a flashlight due to the working environment for the grader at the sawmill measurement station, thus the rolling shutter technique had to be used. The amount of motion blur depends on the end face movement during the time for image capture.

The total time for processing each pixel row is given by the following expression (Anon. *PixeLINK*, 2004)

$$t_r = 7.3 \cdot 10^{-6} + 25 \cdot 10^{-9} \cdot (N_c + 10), \quad (1)$$

where N_c is the number of active pixel columns. The total time for imaging is then

$$t = N_r \cdot t_r, \quad (2)$$

where N_r is the number of pixel rows. Capturing an image of size $N_c = 1536$ and $N_r = 1800$ results in a total time of $t = 82.7$ ms for each image.

If the log moves away from the camera with a speed of v m/s, the total distance the log has moved during image capture is $\Delta s = t \cdot v$ m. At Setra Group Nyby sawmill the conveyor belt is usually moving with a speed of 78–85 m/minute (according to personal communications with J. Olsson, VMF Qbera), which is approximately 1.3–1.4 m/s. In worst case this results in a total movement of the log $\Delta s = 0.12$ m during image capture. However the imaged end face generally does not cover the entire image, thus the time for capturing the part of the image with the end face is shorter than 82 ms. Hence the the distance the log moves during acquisition of the end face is more likely about 0.05 m.

The camera is positioned in a rough environment, where large logs passing on the conveyor can hit both the camera and the lamps, which has happened several times during the year of image acquisition. This means that the camera is not positioned exactly at the same place all the time, resulting in some images being slightly out of focus and thus blurred.

Acquisition was carried out monthly, which resulted in as many as over 4000 images. From this large dataset, a few smaller sets have been created which were used during development and evaluation of methods in this work. A training dataset \mathcal{S}_T , consisting of 229 images was created by picking every 20th of all acquired images and afterwards removing those that were not acquired correctly due to e.g., occlusion by another log. An evaluation set \mathcal{S}_E , was created in the same way, including 223 images. Note that no images where the end face was correctly acquired were removed. A typical image from November 2008 is shown in Figure 15.

The NybyOnline setup was also used for acquisition when evaluating the method for counting annual rings, which is described in Paper VII. There, 70 logs were imaged three times when transported on the conveyor. Before the third image session, the end faces were sawn with a chain saw to produce newly sawn surfaces for comparison. These end faces were also depicted before and after sawing using the Canon camera. Figure 3 shows one of these logs, imaged with the Canon camera.



Figure 10: An image from the Forssjö set. The image is rotated 90 degrees counter-clockwise.



Figure 11: An image from the Ultuna set.



Figure 12: An image from the Nyby set.



Figure 13: An image from the Heby2006 set. The image is rotated 90 degrees counterclockwise.



Figure 14: An image from the Heby2007 set.



Figure 15: An image from the NybyOnline set.

4 Synthetic data

Synthetic end face images can be used for developing and evaluating image analysis methods. An algorithm for generating such images was presented in Paper I, including simulations of wood features, effects from harvesting and storage, and image acquisition. In this chapter the algorithm for simulating two-dimensional (2D) end face images is described together with an extension of the method to simulate three-dimensional (3D) data. For the synthetic images, ground truth can be computed mathematically.

4.1 Background

Simulation of surfaces with wood texture has earlier been performed mainly in the field of computer graphics to create patterns that are realistic and visually appealing. For such applications, the appearance is often the main issue. It is important that the texture is realistic, but what causes the appearance, i.e., tree growth is not necessarily considered. However, if synthetic wood is generated as a tool for image analysis research it is important that the data is as similar as possible to real wood. If so, it can be used in development and evaluation of methods, preferably when the amount of real data is small, or if ground truth is lacking.

There are mainly two ways of applying texture to a surface: *texture mapping* and *procedural texture*. Texture mapping has traditionally been used for 2D images, and procedural textures for 3D volumes. In texture mapping a pattern is fitted to a surface. The pattern is stored as a look-up table, typically an image from which the pattern is picked. With procedural texture the pattern is built up using mathematical functions that are evaluated on every point on a surface or in a volume (Ebert et al., 2003). A mathematical model suitable for the pattern ensures that it behaves as intended and complicated surfaces are thus easier to handle than in the case of textural mapping. Procedural texture was presented in Perlin (1985) and Peachey (1985), and is also called solid texture.

Wood is simulated as solid texture in Peachey (1985), using a rather simple model with bright and dark concentric cylinders to resemble earlywood and latewood, respectively. Another approach is taken in Buchanan (1998) where a growth model is used to simulate wood in 3D. The model mathematically grows the wood from a pith and outwards in the volume. A parameter steers the growth, allowing uneven progress for different directions. Branches are modeled in a similar way as the stem with separate growth directions. Two different colours are chosen for earlywood and latewood, respectively, and the colour of each volume element is computed from interpolation between these.

In Lefebvre and Poulin (2000) a photograph of a wooden surface is used to create a similar synthetic pattern. Parameters like pith position, annual ring frequency, ratio between earlywood and latewood as well as growth directions are extracted from the photograph and inserted in a mathematical model. By evaluating the model in the synthetic volume, the real pattern from the photograph is fitted to a simulated surface.

In Paper I synthetic data was generated in 2D using procedural texture. The aim was to produce images resembling real data, with ground truth and adjustable features that could be used in the evaluation of image analysis methods for real end face images. An effort was made to determine parameters representative for real end faces. Data from a training set of approximately 200 images of pine end faces were used to examine features, both statistically and visually. Apart from the wood also other features were simulated, such as effects from harvesting, damage from storage, camera setup including illumination and the image acquisition itself.

4.2 Generating 2D end face images

To produce realistic data, features such as end face size, presence and size of heartwood, number of visible knots, and number of cracks were analyzed from the training set. The analysis was performed to get a rough estimation of how these features are distributed and was done by the author.

Size was measured as the end face radius from the pith to the border. Since end faces are not completely circular, one direction must be chosen for this measure. Here, the direction with the largest distance from pith to bark was picked by visual inspection. The relative frequency of measured size from the training set samples is shown as a histogram in Figure 16. The shape suggests that it can be described rather well with a normal distribution, a Gaussian function. Such a function, with mean $\mu = 125$ and standard deviation $\sigma = 30$, is plotted in the same figure. The function is used as the distribution from which the radius r is drawn in the simulations.

Presence of heartwood was determined binary, and the measurements showed that heartwood appears in approximately 75% of the logs, hence the probability chosen for creating heartwood in synthetic images was set to 0.75. The size of the heartwood area was determined similarly to the method for determining end face size. The ratio between heartwood size and end face radius is shown in Figure 16. The plotted Gaussian function has mean $\mu = 0.55$ and standard deviation $\sigma = 0.1$ and is the distribution used to determine the heartwood size.

The number of knots and cracks were counted in each image as integer numbers. The result is shown in Figure 16 for the training set samples together with the probability distributions M_k and M_c , used. These numbers were limited to five in the synthetic data to simplify the simulations of them. In some end faces the cracks were short and very many, starting a few cm away from the pith. The exact number of cracks were not counted in these end faces, but was set to five. To be sure not to simulate too many end faces with five cracks because of this, the measured end faces with five cracks were divided between the bins for three, four and five cracks for the synthetic distribution.

The simulations are divided into five different parts according to the following:

- **Wood**

Simulation of wood includes pith position, end face size, presence and size of heartwood and knots, annual ring pattern, bark thickness, and colour. The size of the end face and possible heartwood is measured from the pith, which

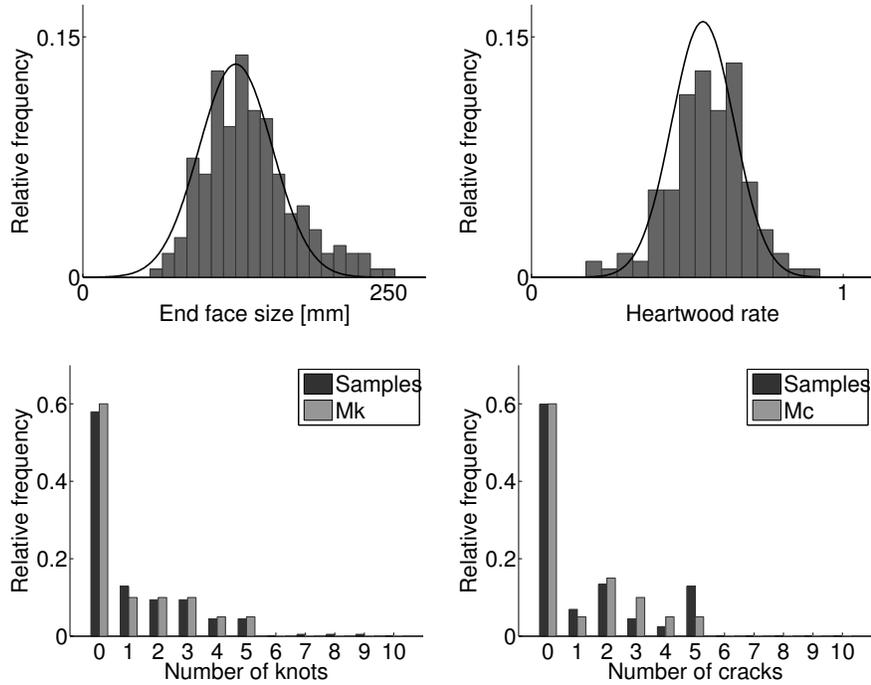


Figure 16: Sample distributions from the training set together with the distributions used for parameters in the synthetic data.

is initially placed in the centre of the image. The knots are defined by a size and a position compared to the pith. The annual rings are simulated using a sine function, $\sin(2\pi fD)$, where D is the signal on which the function is applied. This signal is based on the Euclidean distance to the pith, but to avoid completely circular rings with constant frequency, perturbations are added to it before applying the sine function. By these perturbations, the frequency of the signal will change while the parameter f is held constant. In this way it is possible to compute the ground truth for the number of annual rings, which will be discussed in Section 4.2.1. The signal is manipulated so that the contrast is reduced and the latewood regions thinner than the earlywood. Finally, the bark and colour is simulated.

- **Sawing**

Sawing affects the end face mechanically by creating a rough surface which is here simulated using noise. Uneven sawing can introduce marks or stripes across the end face. The simulation of sawmarks is performed by bump mapping, where the unevenness is created as a matrix of normal directions from the surface (Blinn, 1978). The normal directions are not visible in the image itself, but the effect can be seen when the synthetic end face is illuminated in a later step of the simulations.

- **Storage**

During storage of logs, the wood can dry and cracks occur on the end face. This is simulated here, defining each crack by a length and an angle. The cracks are propagated from the pith and outwards.

- **Camera setup**

The camera setup includes determining the position of two lamps and the camera compared to the end face. In this step the simulated sawmarks will be visible due to a multiplication of the simulated normal map and the directions to the lamps. Except from the marks from sawing, the surface is assumed to be Lambertian, i.e., perfectly diffuse. A geometric distortion is introduced in this step, created from the camera position in relation to the end face.

- **Image acquisition**

The final step is the imaging itself. Here, the end face is resampled to a smaller size and thus lower resolution. Smoothing is performed before resampling to avoid aliasing and at the same time simulate the incoming light spreading in the camera optics. A colour camera is simulated here, and thus Bayer filtering and demosaicking was implemented (Ramanath et al., 2002). Note that simulations of acquiring an intensity image using a colour camera should also use the Bayer filtering for acquisition, and convert to intensity afterwards. If an intensity camera is simulated, the Bayer filtering should be ignored.

The results after different steps of the simulations are shown in Figure 17. The colour noise, which is created in the Bayer filtering is illustrated in Figure 18. A small part of a simulated image with colour noise is shown. The difference between this image and one without Bayer filtering is given for each colour band. The absolute values of the differences are shown as intensity maps, with values between zero (black) and 20 (white). The effect is largest at rapid intensity changes, e.g., cracks, since the Bayer filtering has a smoothing effect on the signals.

The final result after simulation is shown at the top of Figure 19, as an intensity image. This is the image from which the illustration in Figure 17 and Figure 18 are taken. The end face shown at the bottom of Figure 19 is simulated with the same parameters as the end face at the top, but with different perturbing factors. These influence many features so that they vary between the images: annual ring pattern, end face size, bark thickness, heartwood region, and end face colour.

4.2.1 *Ground truth*

Determination of ground truth for the synthetic wood must be considered during simulations. When it comes to the wood features, they are not only affected by the wood simulations themselves, but also by the camera setup and image acquisition. There, the end face distorted due to a projection onto the image plane and resampled to lower resolution. By determining the wood features during simulations, ground truth can be computed from knowledge of how the end face is transformed in the other steps.

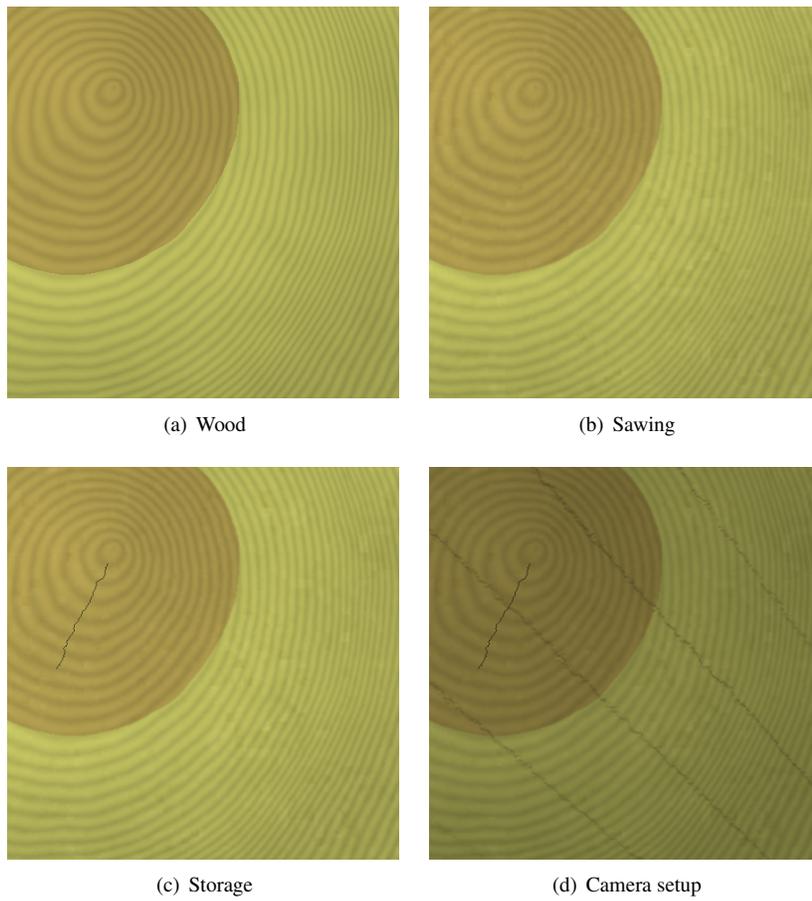


Figure 17: Result after different steps in the simulation.

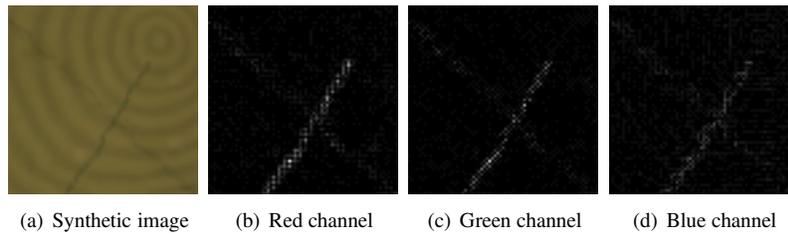


Figure 18: A synthetic image simulated using Bayer filtering shown together with the effect for the red, green and blue colour channels.



Figure 19: Two synthetic images simulated using the same parameters, except for different perturbing factors.

The pith is determined from the beginning as a pixel position. When it comes to knots and cracks the information is naturally stored from the masks creating these features. The number of annual rings can be determined from the frequency f together with D , the perturbed distance function on which the sine function is applied. The period time for the function, and thus the number of pixels for one complete annual ring is determined as $T = 1/f$. For a 1D signal, the number of periods between two points x_1 and x_2 is thus given by $(x_2 - x_1) \cdot f$. Similarly, for two pixels in the same direction from the pith, but at different distance, at positions (x_1, y_1) and (x_2, y_2) the number of rings between them is computed as $(D(x_2, y_2) - D(x_1, y_1)) \cdot f$. Thus storing the distance D and transforming it using the same camera position and image acquisition parameters as for the end face itself, gives a lookup table for D in each pixel of the simulated image.

4.3 Additional experiments - 3D sawlogs

A further development of the method to produce 3D synthetic images, i.e., synthetic logs is shown here. This is performed by adding a third dimension to the previous simulations. Much is similar to the 2D end faces, but a few things must be modified. The pith position, which is a point in the 2D case, is elongated throughout the third dimension. To avoid a pith position constant throughout the log, low frequency noise is added to vary the position. The distance used to simulate annual rings is now computed from the elongated pith. By multiplying the distance with a thinning factor, a log that tapers towards the top of the tree is created. This is taken into account when computing the annual rings, and when determining the heartwood size and the radius of the log.

The largest difference in the simulation is the model for the knots. In the 2D case they were distributed on the end face by choosing a position and the knot size. In 3D, however, the knots at slightly different heights must of course match, therefore it is not possible to place them randomly. Andreu and Rinnhofer (2003) model knots by a spline, but here, for simplicity, the knots are modeled as cones from the pith and outwards. With this model three angles are needed to define each knot: the angle distributing it around the circumference (Θ_k in Paper I), an angle that determines the shape of the cone, and finally an angle that determines how the branch is positioned in relation to the stem growth direction. The number of knots is set to $M_k = 5$ for all whorls. The whorls were placed repeatedly along the log. Examples of simulated logs cut lengthwise are shown in Figure 20. The left and right logs have visible heartwood.

In the 2D simulations, no relation between knot size and their distance to the pith was taken into account. Since branches grow from the pith and outwards, the size should be smaller towards the pith and larger further out towards the bark. This will be the case in the 3D simulations, since a model is used where the branches are simulated as cones. The 3D knot model is thus more realistic and should be used also in the 2D case.



Figure 20: Examples of cross sections of synthetic logs.

Part II
Theory

5 Image analysis

In this work many different image analysis methods were used, which was required to solve the variety of tasks for the real data. This chapter describes a few basic concepts, e.g., thresholding, smoothing, edge detection and mathematical morphology, mostly illustrated here on end face images. Some of the methods are used in this dissertation and others in similar applications. In this work mostly intensity images are handled, thus they will be in focus. However, conversion from colour to intensity must first be performed.

5.1 Pixelwise operations

As mentioned in Section 3.1.3, a colour image consists of 3 colour channels, each represented as an intensity image. Conversion from colour to intensity, I can therefore be performed by choosing one of the colour channels as a representative for the complete image. However, instead combining the different channels into one often results in more image features being conserved after the conversion. There are several ways to do the conversion. A common one is the Y-channel in the YIQ-space,

$$I = 0.2989 \cdot R + 0.5870 \cdot G + 0.1140 \cdot B, \quad (3)$$

where R , G and B represents the red, green and blue colour channels, respectively. This intensity image is related to the way humans interpret colours, with high sensitivity in the green colour band (Sonka et al. (1999)). Another method is to take the mean value of the colour channels,

$$I = \frac{R + G + B}{3}, \quad (4)$$

which is equal to the V-channel in the HSV colour space (Gonzalez and Woods, 2002). Which conversion to use should be determined by the information in the colour image and the application in mind.

The images used in this work were all captured in colour but were converted in two different ways to intensity images. For all datasets except the NybyOnline set Equation 4 was used, chosen by visual inspection. For the NybyOnline dataset, the illumination was different compared to the other sets and the camera, PixeLINK, was also quite different from the others. With this setup it turned out that the blue colour channel had rather low intensity and at the same time was noisy. Therefore this channel was ignored in the conversion from colour to intensity. A modified version of Equation 3 was used, $I = 0.2989 \cdot R + 0.5870 \cdot G$.

5.1.1 Image histogram

Once the intensity image is acquired, either from imaging or, e.g., after conversion from a colour image, its characteristics can be analyzed in various ways. One is by using the image *histogram*, where pixels of equal intensities are binned together to a one dimensional signal. The histogram can be used to analyze image brightness and contrast and it can be a helpful tool in segmentation of the image into object

and background. In Figure 21 the histogram for the green channel in the dandelion image in Figure 6 is plotted. Here, each bin represents one intensity value, resulting in a histogram with 256 bins. The shown histogram has a dominating peak corresponding mostly to background, with values in the range of approximately 50 to 120. Pixels with higher intensities representing mostly the yellow flower and green straws, form a wider, lower peak in the histogram.

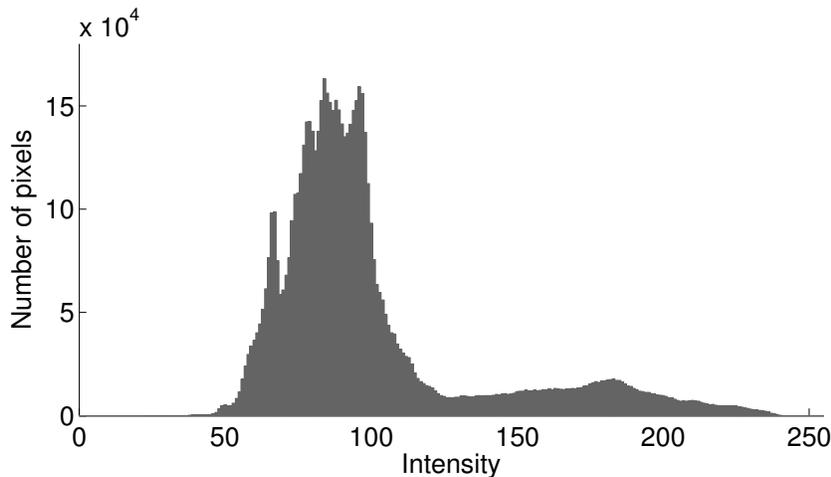


Figure 21: Histogram for the green channel in Figure 6.

5.1.2 Thresholding

The shape of the histogram in Figure 21 suggests that the image can be divided into object (flower) and background using green colour channel intensity. By performing *thresholding*, the image can be converted into a binary image, where 1 corresponds to image pixels with intensity higher than the threshold T and 0 corresponds to pixels with lower intensity,

$$I_T(x, y) = \begin{cases} 0 & \text{if } I(x, y) < T \\ 1 & \text{otherwise.} \end{cases} \quad (5)$$

The threshold T can be chosen manually, automatically or semi-automatically.

Figure 22 shows the dandelion image after thresholding the green colour channel using four different thresholds. The higher the threshold, the more of the image is considered background. In Figure 22(b) thresholding is performed at $T = 128$, which corresponds to a local minimum of the histogram, roughly separating dark background from the bright flower. However, a few straws are also considered object using this threshold.

Global thresholding is often not ideal for separating background and object in this way. It is seldom possible to find a threshold suitable for the entire image. This is illustrated in Figure 23, with an intensity image showing a small part of the

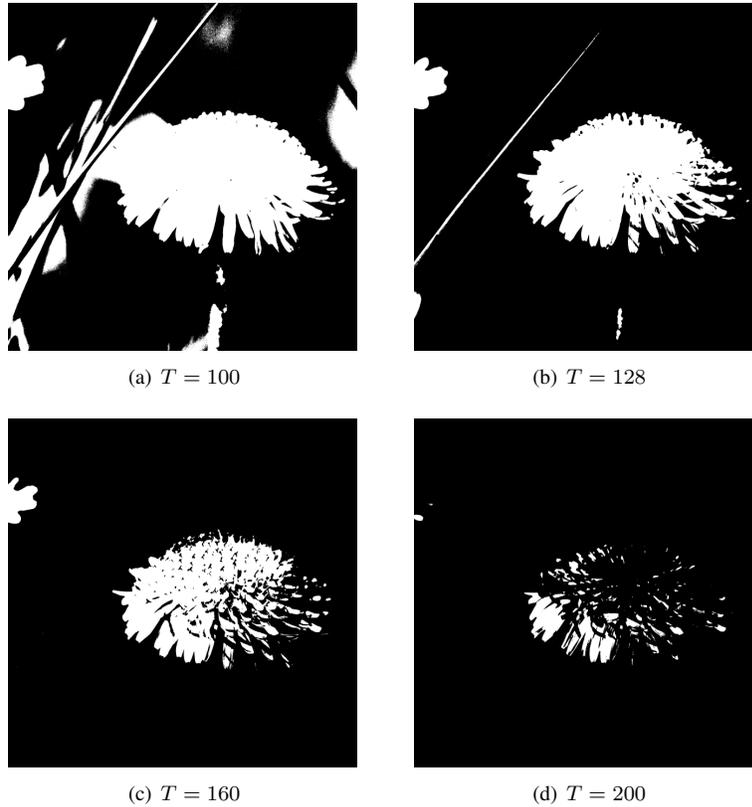


Figure 22: Results after thresholding the green colour channel in Figure 6 for different values of T .

end face in Figure 2, converted to intensity according to Equation 3. The original image is shown together with the result after thresholding at $T = 150$, with the aim to separate the dark latewood from the brighter earlywood. The result shows that a global threshold is not suitable here, since the threshold used is too high in the upper left area and at the same time too low in the lower right area.

Global thresholding is used in several methods for detection of annual rings in wood, e.g., in Wang (1998) for X-ray images of logs. Thresholding a small part of the end face is performed in Chalifour et al. (2001), as a first detection of annual rings. Österberg (2009) concludes that global thresholding is not applicable to separate earlywood from latewood on end face images for several reasons. One is that the colour of both earlywood and latewood varies over the end face and even within one specific ring, giving the same effect that was illustrated in Figure 23. Another reason is that compression wood is often almost as dark as latewood, and thus a complete annual ring with compression wood can be determined as latewood by global thresholding.

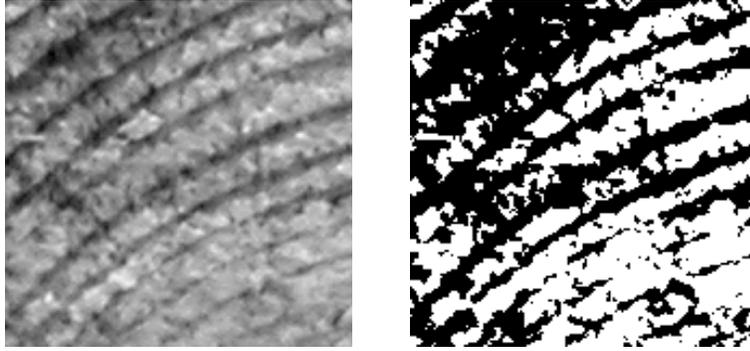


Figure 23: Intensity image (left) thresholded at $T = 150$ (right).

5.2 Operations on pixel neighbourhoods

In histograms and thresholding only the pixel intensity is taken into account, but aside from intensity each pixel is also defined by its spatial position in the image. The position is referred to by the pixel's column and row in the image; the x and y coordinate, respectively. Adjacent pixels are called neighbours and a small area around a pixel can be referred to as the pixel neighbourhood. For a pixel, p , the 4-neighbours are the pixels connected by an edge. By adding the pixels that are connected by a vertex, the 8-neighbours are defined. The 4- and 8-neighbours are shown in Figure 24. The 8-neighbours are equal to the 3×3 neighbourhood of p .

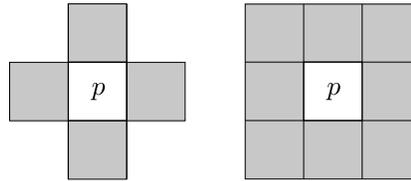


Figure 24: The 4-neighbours (left) and 8-neighbours (right) for a pixel p .

5.2.1 Filtering

The spatial position provides possibilities to perform *spatial filtering* of an image. In spatial filtering, a local neighbourhood of a pixel is considered when mathematical operations are performed. The result of the filtering is determined by the filter mask.

To filter an image I with a mask w is to place the filter mask over each pixel of the image and multiply it pixelwise with the pixels that are covered. The mask has size $m \times n$ pixels where m and n are odd integers and the centre pixel of the mask is the origin, which is placed over each image pixel. The result g of filtering pixel (x, y) is thus given by

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot I(x + s, y + t), \quad (6)$$

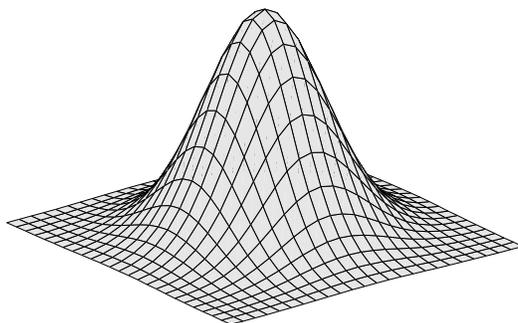


Figure 25: An example of a 2D Gaussian function.

where $a = (m - 1)/2$, $b = (n - 1)/2$, and the centre pixel in w is denoted $(0, 0)$. This procedure is called convolution and the filter mask is the convolution kernel mirrored around its origin. The operation is only defined where the entire filter fits inside the image, i.e., not at the image borders. Filtering can still be performed there, if the image is enlarged, either by mirroring or copying the outermost rows and columns. An image can also be padded by zeros or the mask can be modified at the borders, including only elements which fit inside the image.

A filter with constant values of $w(s, t)$ computes a mean value of a local region. By letting the sum of the weights be one, the intensity range of the image remains approximately unchanged. However, the square shape of the filter introduces square artifacts in the filtered image.

Another typical filter is the *Gaussian filter*, which is a Gaussian function sampled on a grid. A 1D Gaussian function is given by

$$G(x) = \frac{1}{\sqrt{2\pi} \cdot \sigma} e^{-(x-\mu)^2/2\sigma^2}, \quad (7)$$

where μ is the mean value and σ the standard deviation, determining the position and width of the function, respectively. The constant in front of the exponential function sees to that the integral over x is equal to one. Similar to the case with the mean filter, this controls the intensity range of the result. The filter mask size is best determined by the values of the Gaussian function. If the function is not close to zero at the filter mask borders, also this filter will introduce artifacts in the result.

A 2D Gaussian function with equal standard deviation in the two spatial dimensions is shown in Figure 25. The values of the function are plotted as heights. Gaussian filtering in 2D can be performed by two 1D functions, filtering first in one dimension and then in the other. Since the Gaussian filter is separable in this way, it is fast to compute.

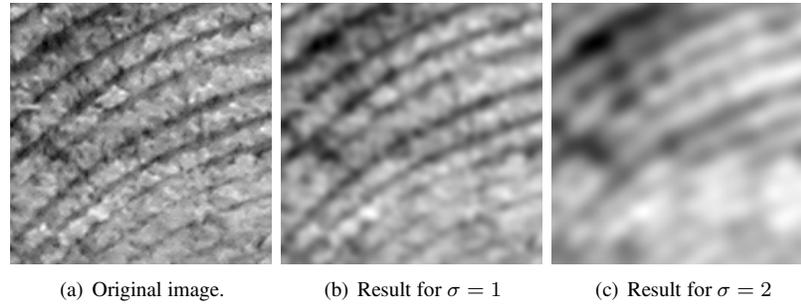


Figure 26: Gaussian filtering using different standard deviation σ .

The Gaussian filter smooths the image, thus reducing high spatial frequencies. The degree of smoothing depends on the standard deviation: the larger the value, the larger is the smoothing effect. This is illustrated in Figure 26, where the small wood image is filtered with two different Gaussian filters. The filter with the larger standard deviation reduces more of the high frequencies. In Chalifour et al. (2001) and Entacher et al. (2007) Gaussian filtering was used as a preprocessing step to remove noise before detecting the pith position and the annual rings.

Another application of spatial filtering is to detect edges in images. One method for edge detection is to estimate the gradient magnitude of each pixel, and threshold the result. The gradient is an approximation of the derivative of the signal sampled in the image pixels. The Sobel operator is a typical tool for computing the gradient, which consists of two 3×3 filter masks that detect edges in the horizontal and vertical coordinates, respectively. Each of these masks are shown in Figure 27.

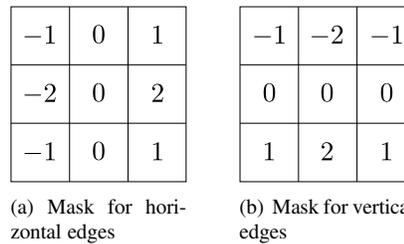


Figure 27: Filter masks for the Sobel operator.

Results after filtering with the Sobel operator are shown in Figure 28 for a sanded end face and an untreated end face. The two images are taken from the same wood disc, but from different sides. This is a part of an end face from the Forssjö set, which was sanded on one side, and completely untreated on the other side. The images are taken from approximately the same region of the disc, but on different sides. The original images are shown together with the computed answers from the Sobel operator in the horizontal direction, and the result after thresholding the gradient images at $T = 30$. By thresholding the gradient image in this way, the edges can be detected.

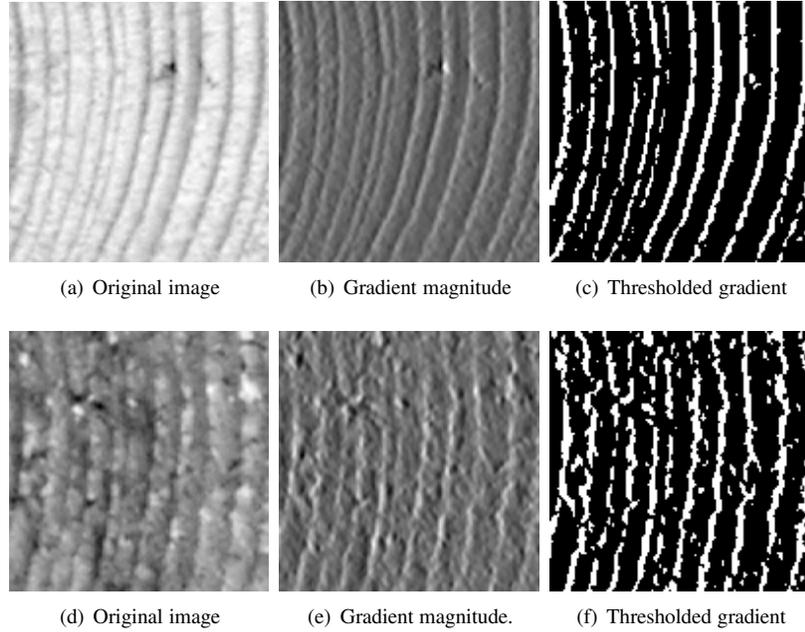


Figure 28: Result after edge detection with the Sobel operator for a sanded end face (upper) and an untreated end face (lower).

5.2.2 Local orientation

The region around a pixel can be categorized in different ways, depending on the pixel intensities and the pattern they create. If the pixel intensities for a small part of an image form a pattern that varies in one direction only, it is called simple. A simple signal can be described as

$$f(\mathbf{x}) = h(\mathbf{k}^T \cdot \mathbf{x}), \quad (8)$$

where h is a function in one dimension, \mathbf{k} is a 2D vector of unit length, and $\mathbf{x} = [x_1 \ x_2]^T$, with x_1 and x_2 as the horizontal and vertical coordinate, respectively. The signal $f(\mathbf{x}) = \sin(\mathbf{k}^T \cdot \mathbf{x})$ is plotted in Figure 29(a), where $\mathbf{k} = [0.86 \ 0.51]^T$ is the orientation along which the signal varies. The local orientation is marked with a line across the pattern. For the perpendicular direction, the pixel values are constant. A small image of parallel annual rings is an image that is approximately simple but corrupted with noise, which is illustrated in Figure 29(b). The orientation is similar to the synthetic image. The signal is approximately constant in the direction perpendicular to the local orientation. In this work the local orientation is used to estimate the pith position in Paper V. The local orientation of annual rings have been estimated using different techniques, e.g., with Quadrature filters (Granlund and Knutsson, 1995) or Laplacian pyramids (Bigün, 1992), which was both used in Paper V, or with the fast Fourier transform (Cooley and Tukey, 1965) as in Hanning et al. (2003) and Österberg (2009).

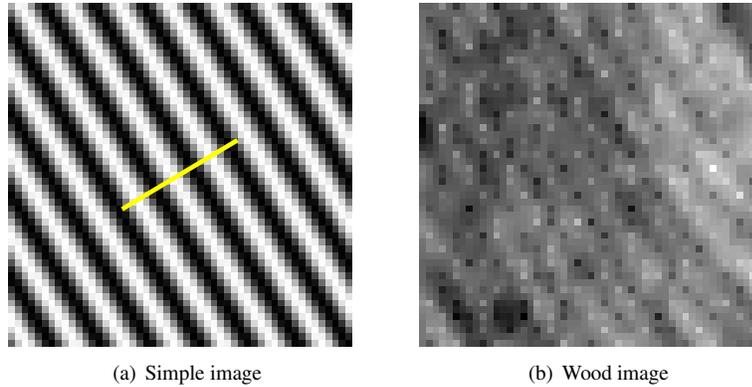


Figure 29: Example of a synthetic, simple image and an image with annual rings which can be approximated as a simple image. The line across the synthetic pattern is the local orientation.

5.2.3 Binary mathematical morphology

Another family of operations regarding spatial relations between pixels is mathematical morphology, where set operations are applied to images (Serra, 1982). Two basic operations are dilation and erosion. Dilation is defined as an intersection between two sets

$$\delta_B(X) = \{x | (\hat{B})_x \cap X \neq \emptyset\}, \quad (9)$$

where \hat{B} denotes the reflection of B in its origin and $(\cdot)_x$ is the translation by a vector x . Thus, $\delta_B(X)$ is the set of all pixels such that X and the reflection of B overlap in at least one pixel when B is shifted over the image. The set B is the *structuring element* (SE), which is often a small set of only a few pixels. If the SE is symmetric around its origin, $B = \hat{B}$. Dilation expands a set, while erosion shrinks it. Erosion is defined as

$$\varepsilon_B(X) = \{x | (B)_x \subseteq X\}, \quad (10)$$

which is the set of all pixels so that B is contained in X .

The result of the morphological operation depends on the shape of the SE which is illustrated in Figure 30. The sets are shown as black pixels. The upper image shows the original set and under it three different structuring elements (a square, a line, and a cross) are shown together with the results after dilation. The origins of the SEs are shown as dots in the centres of them. Dilation with all these SEs fill holes in the letters B and A and at the same time expand the letters to larger objects with different outer shape.

By applying dilation followed by erosion using the same SE *closing* is performed. With this operation small holes in a set are filled. In the same manner erosion can be followed by dilation, called *opening*, with the result that small objects are removed. An example of closing is illustrated in Figure 31. This is the result after performing dilation followed by erosion on the original image seen upmost

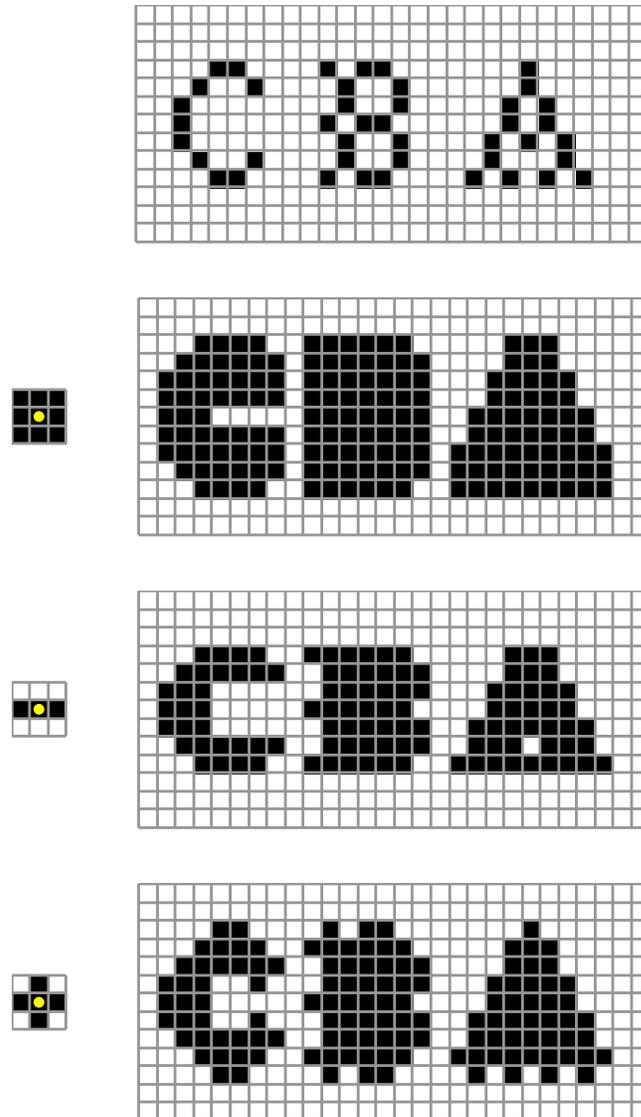


Figure 30: Dilation of a set (uppermost image) using three different SEs (left) giving three different results (right).

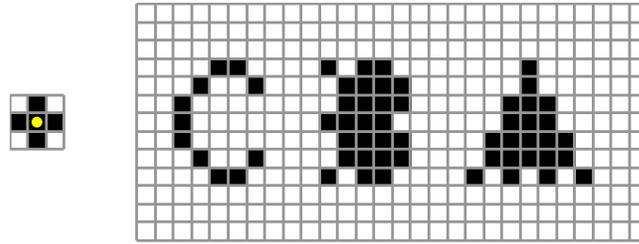


Figure 31: Closing performed on the uppermost set in Figure 30 using the cross SE.

in Figure 30, using the cross SE. The difference between closing and dilation (bottom row in Figure 30) is that the holes are still filled, but the size of the object has shrunk again, and the outer contour is more similar to the original set.

Binary morphological closing is used in Wang (1998) to improve the result after detecting annual rings by thresholding X-ray images of logs. Mathematical morphology can also be performed on intensity images, which will not be described here, but was used by Soille and Misson (2001) to measure annual ring area in prepared end faces acquired using a microscope.

By letting the SE vary in shape and size with the spatial coordinates of the image, *spatially variant* mathematical morphology is performed. In Heijmans and Ronse (1990) polar coordinates are used to change the size of the SE according to the radial coordinate. The SE can also be adapted to the image data, e.g., for connecting broken edges (Shih and Cheng, 2004).

6 Distance transforms

A distance transformation changes an image, binary or grey scale, into a distance transform or distance map. There, each pixel value corresponds to the distance from the pixel to the object from which the distance was computed. There are several ways to define distance in an image, using different metrics, and the choice depends on the application in mind. In this chapter distance transforms are described including the polar distance transform which was first introduced in Paper II and used also in Paper III and Paper IV. The polar distance transform is an important tool for outlining annual rings in this work.

In image analysis distance is often computed by propagating distance values through the image. This can be performed using several methods with different results. Ideally, a distance transform should be rotationally invariant, so that the rotation of the grid affects the result as little as possible. Different techniques and metrics gives more or less rotational variance. Here, two algorithms for computing the distance transform are described: *chamfering* and *fast marching*. The *polar distance transform* was computed using both of them.

By taking pixel intensities into account in the computations the distance is *grey weighted*. Then distance between pixels at specific positions can be different depending on the intensities in the image. The *grey weighted polar distance transform* is also computed in this work using both chamfering and fast marching.

6.1 Computations using chamfering

In chamfering a path based distance is computed, thus the distance between two pixels is defined by the shortest path between them. The path is built up by a sequence of steps, each taken in a local neighbourhood. If a path between the pixels p and q is given by $\mathcal{P} = \langle p = p_0, p_1, p_2, \dots, p_n = q \rangle$ then the cost is defined as the sum of the distances for each step, $\mathcal{C}(\mathcal{P}) = d(p_0, p_1) + d(p_1, p_2) + \dots + d(p_{n-1}, p_n)$. The total distance is defined by the path that gives the lowest cost, thus

$$d(p, q) = \min_{\mathcal{P} \in \Pi} \mathcal{C}(\mathcal{P}), \quad (11)$$

where Π is the set of all paths between p and q . The distance between two neighbouring pixels is defined by the length of the step, a *weight*, thus a weighted distance is computed (Borgefors, 1986). Different weights can be used for different types of neighbours, e.g., edge neighbours and vertex neighbours.

Examples of masks with weights in two different neighbourhood sizes are shown in Figure 32. The letters denote the different types of steps that occur, i.e., steps to edge neighbours (type a), vertex neighbours (type b), and knight neighbours (type c). The empty positions in the larger neighbourhood correspond to pixels that can be reached from two shorter steps of the same kind, and are therefore redundant.

It was shown by Rosenfeld and Pfaltz (1966) that a path based distance can be computed using sequential operations and that two raster scans over the image is enough. Initially, the distance is set to zero for the pixels that belong to the object

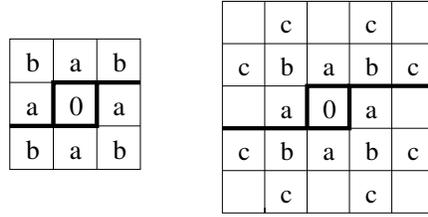


Figure 32: Masks for computation of the weighted distance for 3×3 and 5×5 neighbourhoods. The thick lines indicate where the masks are split in chamfering. The origin is in the centre of each mask and is included in both parts.

from which the distance is computed, and to infinity for the pixels in the background. The image is scanned from the upper left corner using the upper left part of the mask (forward scan) followed by a scan from the lower right corner using the lower right part of the mask (backward scan). This is called chamfering. The thick lines in Figure 32 marks where the masks are divided into upper and lower part, with the centre element included in both parts.

An illustration of a forward scan using a 3×3 mask is shown in Figure 33. Only the upper left part of the mask is shown, which is the part considered in the forward scan. The mask origin is marked with a dot. The grey part of the image marks visited pixels, so far. The computations are only valid where the entire mask fits inside the image. Computations for other pixels can either be ignored and set to infinity, or the image can be padded with extra rows and columns so that the mask can fit inside the original image. In Figure 33, border pixels are ignored.

The distance for a pixel, p , is computed by propagating distance values from the neighbouring pixels in the distance map, D adding the weight for the local step. The result with the lowest value is the distance in p ,

$$D(p) = \min_{v \in \text{upper}} (D(p + v) + w(v)), \quad (12)$$

where v is a vector from p to a pixel in the mask and $w(v)$ is the corresponding weight in the mask. After this scan, low distance values have been propagated from the upper left part of the image. To compute the final result the backward scan must be performed,

$$D(p) = \min_{v \in \text{lower}} (D(p + v) + w(v)). \quad (13)$$

The size of the mask together with the weights, determine the path based distance, and thus the degree of rotational invariance. The result often improves with a larger mask, but results can be fair even with a small neighbourhood if the weights are chosen properly. Borgefors (1986) recommended the integer weights $a = 3$, $b = 4$, in the 3×3 neighbourhood. Examples of three different sets of weights that have different properties when it comes to rotational invariance are shown in Figure 34. The weights $a = 1$, $b = 1$ results in the so called chessboard distance, and $a = 1$, $b = \infty$ (i.e., the vertex neighbours are not used) results in the cityblock distance.

Since the computed distance is path based, i.e., propagated from pixel to pixel, it is possible to keep track of which neighbour propagates a distance value to a pixel. A path from the object to a pixel can therefore be determined from all the local steps on the way. Often, several paths have equal cost, and no unique path exists. Instead, there are several shortest paths.

6.1.1 Grey weighted distance

In grey weighted distance, pixel intensity is taken into account in distance computations. It was first presented in Rutovitz (1968), where the distance is defined as the sum of the intensity values of the pixels along a path. No weights were used for propagating the distance values, only image intensity is taken into account.

Levi and Montanari (1970) presented a version of the grey-weighted distance transform (GWDT) in which the length of the arc connecting two pixels is weighted with the mean of the intensity values of the two pixels in the step,

$$c(p, q) = \frac{g(p) + g(q)}{2} \cdot d(p, q), \quad (14)$$

where $g(p)$ is the intensity value in pixel p and $c(p, q)$ is the total cost for the local step. Paths along low intensities will have low cost compared to paths covering pixels with high intensities. Levi and Montanari (1970) also concluded that two scans are not enough for these computations, thus iterations must continue until the distance map has converged.

Another approach was taken by Toivanen (1996) in the weighted distance transform on curved spaces (WDTOCS), where the difference in grey values between two pixels in a step is added to the length of the step. This method thus prefers paths where the intensity is constant. A scaling parameter is included to determine to what extent the image intensities influence the distance.

Fouard and Gedda (2006) concluded that the main difference in appearance between GWDT and WDTOCS are that the first one smoothes the cost function, while the latter rather enhances the differences. The fact that the WDTOCS adds the intensity to the distance also makes it more sensitive to scaling of the cost image.

6.2 Computations using fast marching

Another method for distance computations is the fast marching method, where a solution to the Eikonal equation is approximated in each pixel (Sethian, 1999). Distance is propagated outwards a wave front from a set of pixels. The speed of the propagation is determined by a speed function, F and the distance can be measured as the arrival time, T when the wave front reaches the pixel. The arrival time and speed function are related through the Eikonal equation, $|\nabla T|F = 1$, or

$$\left(\frac{\partial T(a, b)}{\partial x}\right)^2 + \left(\frac{\partial T(a, b)}{\partial y}\right)^2 = \frac{1}{F(a, b)^2}, \quad (15)$$

for a point (a, b) . With a constant value of $F = 1$ the distance is computed on a constant cost image, taking only the spatial positions of the pixels into account.

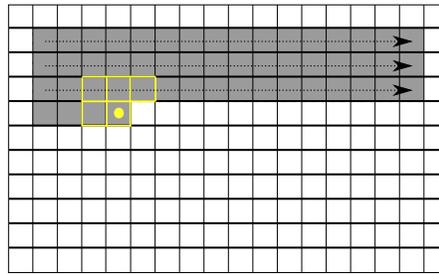


Figure 33: Illustration of the forward scan in chamfering. Grey pixels mark the positions where the mask has already been placed. The mask is shown in yellow.

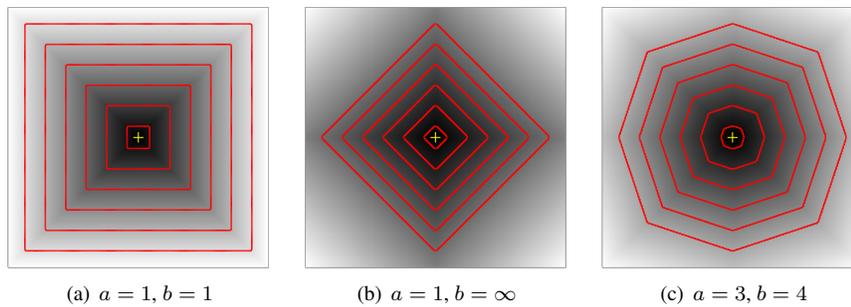


Figure 34: Distance maps shown as intensity images with superimposed level curves for three different sets of weights. The distance is computed from the centre of each image (+).

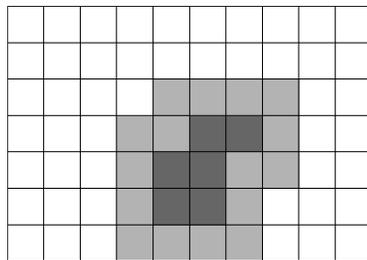


Figure 35: An example of labels of the pixels in fast marching: frozen (dark), narrow (grey), and far (white).

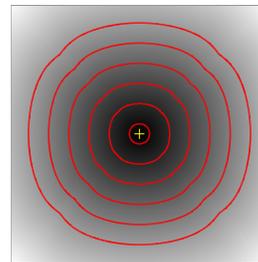


Figure 36: A distance map shown as an intensity image with superimposed level curves. The distance is computed from the image centre (+).

The front propagates from small distance values to larger, updating the new distance values by approximating the solution of the Eikonal equation in each point with approximations of local derivatives. With this method information from one or more pixels are combined when propagating the distance.

Initially, all pixels in the object from which the distance is computed is set to zero and marked as *frozen*. Frozen pixels are set and cannot be changed in the further computations. In the next step the pixels closest to the frozen ones are considered. This is the *narrow band*, and consists of all pixels that are possible for updating, i.e., all pixels neighbouring the frozen pixels. All other pixels are *far away* and are therefore not considered at the moment. For each of the narrow pixels a distance value is computed using the frozen pixels. The narrow pixel with the lowest value (shortest distance) is the next pixel to which the front will propagate, and this pixel is moved to the frozen set. Now, the narrow pixels need to be updated to include neighbouring pixels also to the last frozen pixel and distances must be updated. The narrow pixel with the lowest value is again frozen and the narrow band is updated. Computations are iterated like this, and finished when all pixels are frozen.

In the original method the 4-neighbours of each pixel are considered for possible updating, which results in large errors along the diagonal directions (Bærentzen, 2001). In Hassouna and Farag (2007) the error is reduced by taking into account the 8-neighbours of each pixel, i.e., diagonal neighbours. An illustration of frozen, narrow, and far pixels is shown in Figure 35 using 8-neighbours for approximating the derivatives. A distance map computed using fast marching is shown in Figure 36 where the distance is computed from the centre of the image.

6.2.1 Grey weighted distance

To compute the grey weighted distance with fast marching requires no modification of the method. By letting the speed function, F in Equation 15 vary in the image, the speed of the wave front propagation will also vary. Here we will denote by cost function, the inverse of the speed function. In this way, the cost have similar effect on the distance as the intensity image in the grey weighted distance using chamfering. A high value of the cost function will thus give a large value of $|\nabla T|$ in the Eikonal equation, meaning the arrival time changes a lot between two points, i.e., the wave is propagating slowly. Pixels with low intensities in the cost image will be pixels where the wave front propagates fast and high cost will result in slow propagation.

An application for grey weighted distance transform is shown in Figure 37, where an area of rot has been segmented. Here, the cost image is computed as $c(x, y) = (I_S(x, y) - 110)^2$, where I_S is the original intensity image smoothed with a Gaussian filter of standard deviation 2. In the cost image, the rot area has low intensity values and thus low cost. The distance is computed from the $+$, which is manually placed in the approximate centre of the area. By thresholding the image at $T = 15$, the boundary of the region is determined.

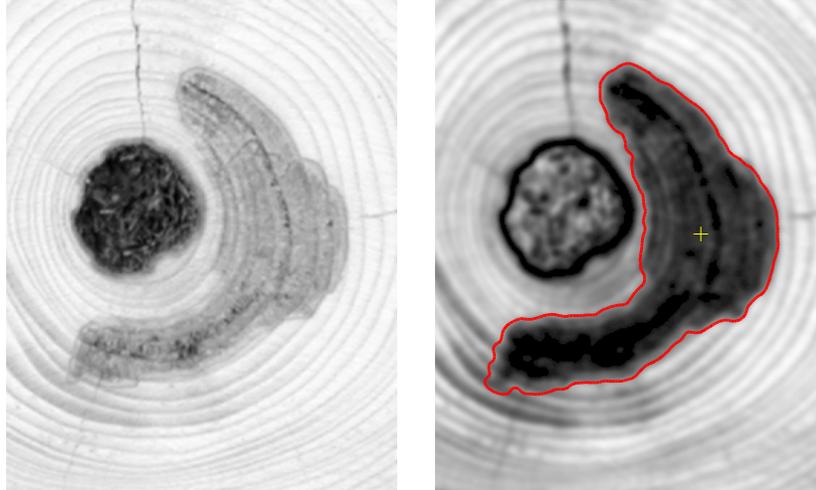


Figure 37: Original image (left) and the cost image with segmented rot area outlined in red (right). The pixel from which the distance was computed is marked with +.

6.3 The polar distance transform

When computing distance with the weighted distance transform or the fast marching method on a constant cost image, the shortest distance between two pixels is an approximation of a straight line. However, in some cases other paths are preferred. With the polar distance transform (PDT) presented in Paper II, the shortest distance between two pixels is influenced by their positions compared to an image origin. For two pixels on the same distance from the origin the shortest path between them is defined as a segment of a circle, with the centre of the circle in the image origin. The angle, θ and radius, r can be computed in polar coordinates for each pixel and by letting the cost for a step depend on the change in radius and angle, the distance can be forced to propagate faster in, e.g., the angular direction than in the radial.

6.3.1 Computations using chamfering

In Paper II the PDT was computed using chamfering. The total weight for a step from pixel p to pixel $p + v$ is given by

$$w(p, v) = \sqrt{(w_r \cdot \Delta r(p, v))^2 + (w_\theta \cdot \Delta \theta(p, v))^2}, \quad (16)$$

where w_r and w_θ are weights assigned to the radial and the angular components of the steps, respectively, and $\Delta r(p, v)$ and $\Delta \theta(p, v)$ are the changes in r and θ for the step from p to $p + v$. The weights in the mask differ for all pixels, since they depend on the pixel's position in the image. In Paper II the PDT was computed using chamfering with masks of size 5×5 pixels due to improved results for rotational invariance compared to the 3×3 neighbourhood. For the polar distance transform, the complete 5×5 mask must be taken into account in the computations, since it cannot be assumed that the weights in certain positions are redundant, as in the weighted distance transform described in Section 6.1.

The choice of the weights w_r and w_θ determines how steps in the radial direction are punished compared to steps in the angular direction. In a digital image, the radius is measured in pixels, and along a few directions it is possible to take steps completely along the radial directions, e.g., on the horizontal or vertical axes, one step towards or away from the origin. For such a step $\Delta r(p, v) = 1$ and $\Delta\theta = 0$, thus $w = w_r$. By letting $w_r = 1$ a step changing the radial component with one pixel, will have weight 1. Furthermore, another step having $\Delta r(p, v) = 0$ would give a total weight $w = w_\theta \cdot \Delta\theta$. For equal propagation in radial and angular direction, this total weight should also be 1, however $\Delta\theta$ for a step depends on the radius. By comparing this with the length of a circular segment, given by $s = r \cdot \theta$, setting $s = 1$, the total weight for this step can be seen as the length of a circular segment, where w_θ is the radius. For equal weights in the two directions weights $w_r = 1$ and $w_\theta = r$ can be used. By increasing w_r (or decreasing w_θ) propagation in the radial direction is punished compared to the angular.

Examples of masks are shown in Figure 38 for a few different weights. The pixel p is positioned at $r = 100$ and $\theta = \pi/6$. Weights $w_r = 1$ and $w_\theta = 100$ yields a mask which is rather symmetric around the pixel p , and thus propagation is similar in all directions. For higher values of w_r the propagation along the radial direction is more expensive, which is illustrated in Figure 38(b) and Figure 38(c). Each step will, however, also change the radial component. Due to this all weights are affected when w_r is increased.

The result after computing the PDT for a few different set of weights are illustrated in Figure 39. The distance maps are plotted as intensity images together with level curves to visualize the result. Weights $w_r = 1$, $w_\theta = R$, where R is the radius for the pixel from which the distance is computed, correspond to similar propagation in all directions. When w_r is larger, propagation to pixels where the radius is changing is more expensive than propagation where the angle is changing, thus the distance propagates faster along the constant radius. This effect is larger for larger w_r . The position of the object compared to the origin is equal to the pixel considered in Figure 38, i.e., $r = 100$ and $\theta = \pi/6$.

The PDT can also be computed as a grey weighted polar distance transform (GWPDT) as was described in Paper II, where grey weighted distance was computed using a 5×5 mask. With a 5×5 neighbourhood, Equation 14 must be modified slightly so that not only the intensity values for pixel p and q are taken into account, but also the intensity for other pixels covered in the step. Otherwise there is a risk that pixels with high intensities are covered in the step, although the cost is not taken into account, e.g., if a knight step is taken. This was also concluded by Ikonen (2006). In Paper II the cost for each step longer than one pixel was taken as the mean value of the intensities in all covered pixels, and the influence from each pixel was weighted by the length of the step covering it.

6.3.2 Computations using fast marching

In Paper III the fast marching method was used to compute the polar distance transform. A local coordinate system with unit vectors \mathbf{v}_1 and \mathbf{v}_2 , corresponding to the radial and angular direction respectively, is placed in the considered pixel. Multiple

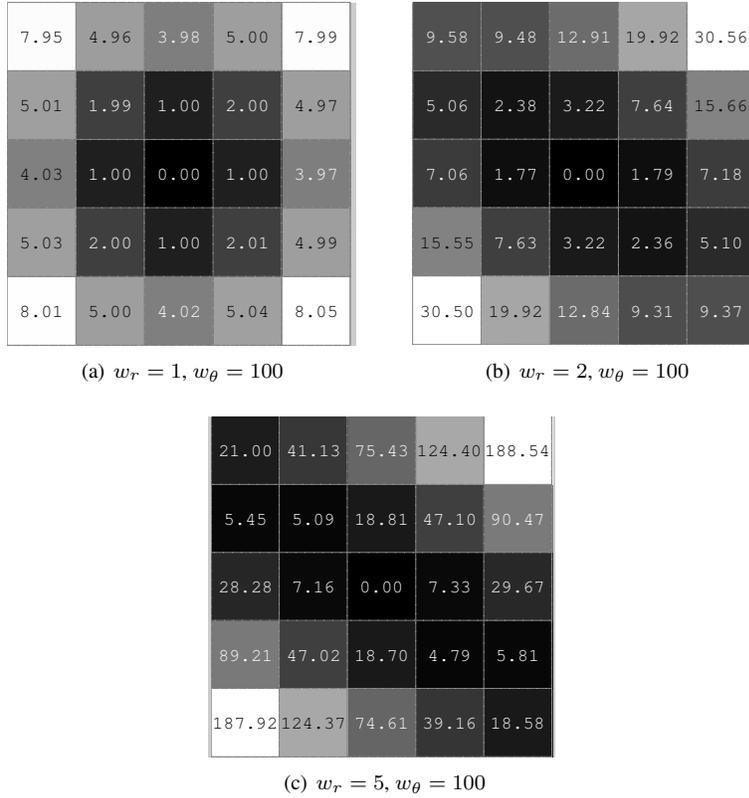


Figure 38: Masks for different sets of weights for a pixel with $r = 100$ and $\theta = \pi/6$.

stencils are used to approximate the solution (Hassouna and Farag, 2007), i.e., the 8-neighbours are taken into account. A weighted version of the Eikonal equation was used

$$\left(\frac{1}{c_1}D_{\mathbf{v}_1}T(a,b)\right)^2 + \left(\frac{1}{c_2}D_{\mathbf{v}_2}T(a,b)\right)^2 = \frac{1}{F(a,b)^2}, \quad (17)$$

where c_1 and c_2 are weights corresponding to the directions \mathbf{v}_1 and \mathbf{v}_2 , respectively and $D_{\mathbf{v}_1}$ denotes the approximation of the derivative along vector \mathbf{v}_1 . The geometry is illustrated in Figure 40.

When computing the PDT with fast marching, equal weights c_1 and c_2 gives similar propagation in the radial and angular directions. By increasing the weight in the radial direction, c_1 , propagation in that direction is punished and thus slower.

The results after computations with a few different weights are shown in Figure 41 where the distance maps are shown together with level curves. The level curves are plotted at the same levels as the distance maps in Figure 39. With the fast marching method, only the 3×3 neighbourhood is taken into account, which is a difference compared to the chamfering computations.

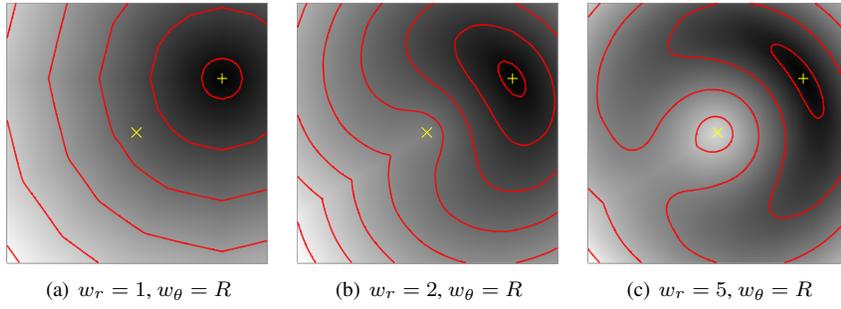


Figure 39: Distance maps for a few different weights computed by chamfering. The distance maps are shown in grey scale together with level curves. The image origins are marked (\times) as well as the object from which the distances were computed ($+$).

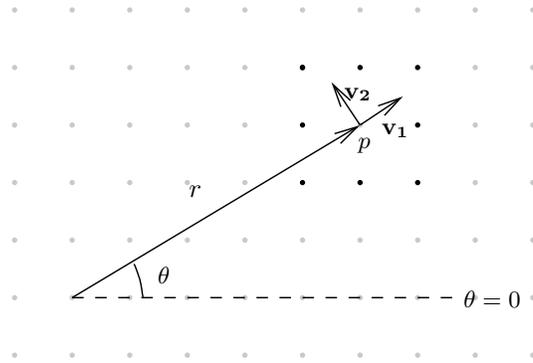


Figure 40: The geometry for the polar distance transform on a digital grid. A local Cartesian coordinate system is placed in the pixel p . The black points represent the 8-neighbours of p .

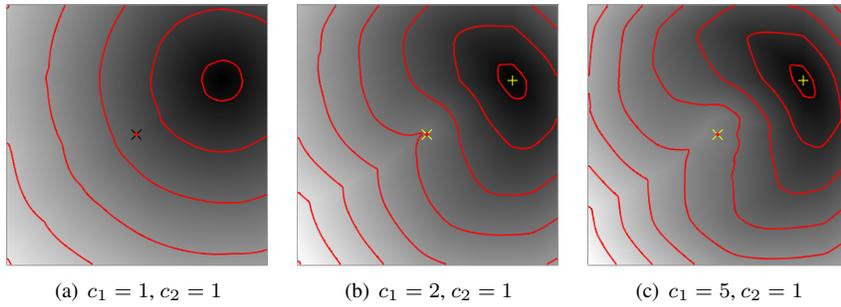


Figure 41: Distance maps for a few different weights computed by fast marching. The distance maps are shown in grey scale together with level curves. The image origins are marked (\times) as well as the object from which the distances were computed ($+$).

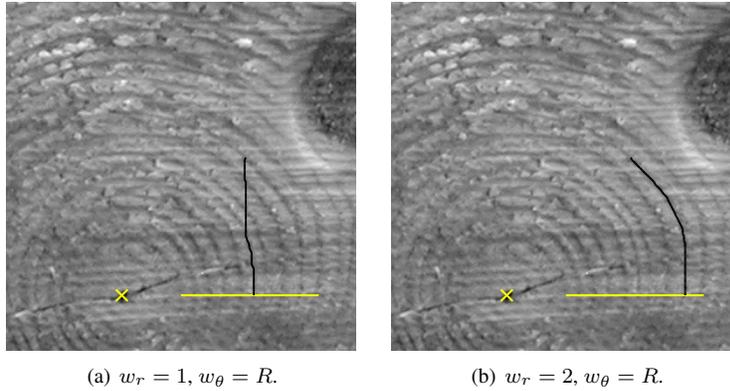
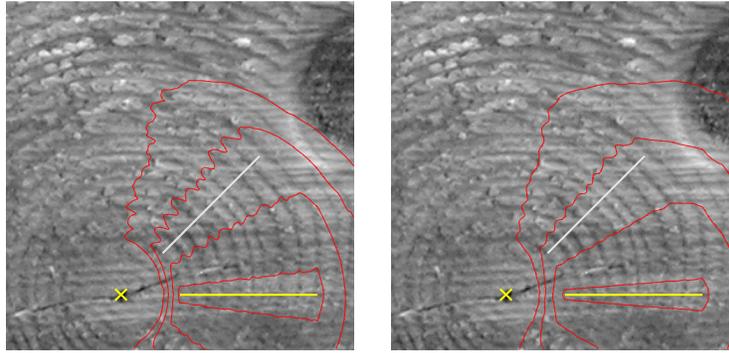


Figure 42: The grey weighted polar distance transform applied to an image with annual ring pattern. The cost image is shown together with the origin (\times), the object from which the distance was computed (line) and the shortest path from a pixel in the image to the nearest pixel in the object (black).

6.4 Applications using the polar distance transform

The GWPDT computed with chamfering can be used for outlining annual rings, as was illustrated in Paper II. The effect of different weights in this application is shown in Figure 42, where the GWPDT was applied to a cost image with an annual ring pattern. The distance is computed from the horizontal line, with image origin in the pith. The shortest path from a pixel in the image to the object is shown as a black path. Results are shown for two different set of weights. With the weights punishing change in radial direction the result is that the shortest path keep to an approximately circular path that also have a low cost, i.e., the latewood region of an annual ring.

Another application is shown in Figure 43, where the GWPDT was computed from the horizontal line, similar to the previous example using both chamfering and fast marching. The cost image is the end face image shown, and level curves for the computed distances are plotted in red. The distance was examined for both transforms along the diagonal white line. The distance values along these lines are plotted in Figure 44 for the original (cost) image, the result after fast marching and the result after chamfering. The signals have been scaled to fit in the same plot without overlapping. The GWPDT using both chamfering and fast marching results in signals that are smoother compared to the original. GWPDT using chamfering has been used as an important step towards counting the number of annual rings in Paper VI and Paper VII.



(a) Computations using chamfering with weights $w_r = 5/R$ and $w_\theta = 1$.

(b) Computations using fast marching with weights $c_1 = 5$ and $c_2 = 1/R$.

Figure 43: GWPDT applied to an image with annual ring pattern. The cost image is shown together with the origin (\times) and the object from which the GWPDT was computed (horizontal line). The distance is shown with red level curves. The white lines marks the pixels for which the distance is analyzed further in Figure 44.

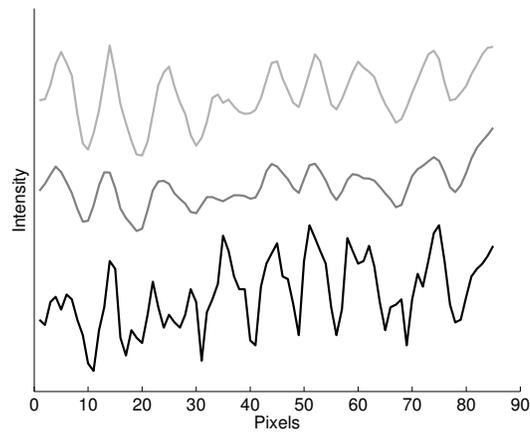


Figure 44: Signals examined along the white lines in Figure 43 for the cost image (black), after fast marching (grey) and after chamfering (light grey).

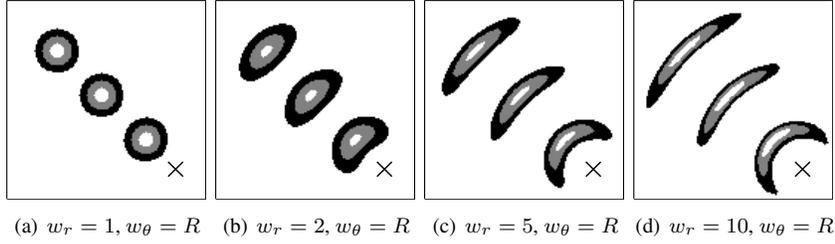


Figure 45: Structuring elements for different weights shown for three different thresholds for each combination of weights. The \times marks the origin. See text for further description.

6.4.1 Binary mathematical morphology using distance transforms

Mathematical morphology can be performed using distance transforms (Soille (1999)). Erosion and dilation of the set A by thresholding a distance map are defined as

$$\varepsilon_D^T(X) = \{x \in X \mid D(x) > T\}, \quad (18)$$

$$\delta_D^T(X) = (\varepsilon(X^C))^C, \quad (19)$$

respectively, where T is the threshold, X^C denotes the complement of the set X and D is a distance transform

$$D(x) = \min_{y \in X^C} d(x, y). \quad (20)$$

For a weighted distance transform, the weights define the shape of the SE. Computing the weighted distance transform using weights $a = 1$ and $b = \infty$, according to the notation in Figure 32 and thresholding the distance map at $T = 1$, would give the same result as in Figure 30 for the cross shaped structuring element. Morphology computed from a distance transform has been used by Tsui (1996) to perform curve thickening of detected annual ring segments in CT images of wood.

By using the PDT in morphology, the shape of the SE will vary over the image, i.e., spatially variant morphology can be performed. In Paper IV it was shown that the PDT is a metric and therefore can be used to compute dilation and erosion. Since the distance propagates differently according to the spatial positions in the image, the SEs will differ, depending on the pixel position. By choosing different weights it is also possible to influence the shape of the SE. Figure 45 shows SEs for different weights at three different thresholds for each set of weights. For each image, one set of weights is used. The structuring elements are shown at three different radii from the origin. For each image, three different thresholds are performed with the resulting SEs in different colours (white, grey, and black). An application was shown in Paper IV, where detected segments of annual rings could be connected by performing dilation with the PDT.

Part III
Applications

7 Preprocessing

An important part of image analysis is preprocessing. This can include noise reduction, contrast enhancement, or other methods that enhance the interesting features in the image. In this chapter the preprocessing required before analysing the end face images is described, including determination of pixel size for all datasets. For convenience, due to the large number of data, an automatic method for cropping the images in the NybyOnline dataset has been developed. The method has previously not been presented, but it has been used for this dataset. The preprocessing also includes an automatic method to reduce marks from uneven sawing, which was presented in Paper VII. This method can be applied to any end face image.

7.1 Determining pixel size

When computing the number of annual rings on end faces as a grading feature, the measurement area is defined as a region of specific size, as described in Section 2.4. To know what this size correspond to in the image, the pixel size must be known. This has been computed in two different ways in this work, depending on how the end faces were imaged.

7.1.1 Images from the Nikon and Canon cameras

For all datasets acquired using the Nikon or Canon cameras, a small reference was imaged along with every end face, as described in Section 3.2. The true size of the reference and the number of pixels it covers in the image are used to determine the pixel size for the entire image. This simple method assumes that the end face and reference are positioned at equal distance from the camera and that each pixel covers a square on the end face, i.e., there is no perspective distortion. This is not always quite the case, but the method is easy to use and gives an estimation of the pixel size.

Measuring the reference size in the image is performed with manual input. However, marking exact corners of the reference is difficult, and sometimes one corner is occluded. Instead the user marks two points on each side of the reference, i.e., eight points in total. By drawing straight lines through the points on each side, the reference is outlined. The intersections between the lines are the estimated corner points. The distance between the corners is compared to the true size of the reference, which is 53×30 mm. Pixel size is computed as the ratio between the size in mm and pixels.

A small part of an image from the Heby2006 set is shown in Figure 46, with the outlined reference. The measures are listed in Table 5, starting with the left side, going counterclockwise. Different sides of the reference give slightly different pixel size. The reason can be both that the reference is rotated compared to image plane so that perspective distortion occurs, and that the reference it is not completely flat. The pixel size is chosen as the lowest value of the four sizes computed, since it is assumed to be least affected by such perturbations.

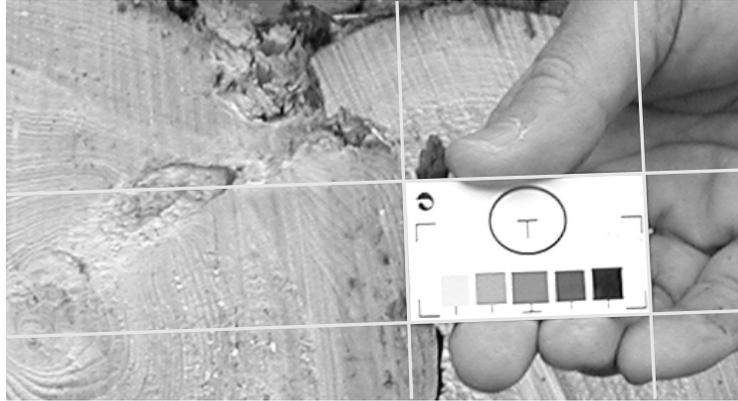


Figure 46: The reference object for an image in the Heby2006 set. The line intersections are the estimated corner positions.

Table 5: Measures for the reference in Figure 46.

length [mm]	length [pixels]	pixel size [mm/pixel]
30	99.7	0.30
53	170.5	0.31
30	99.4	0.30
53	169.2	0.31

7.1.2 Images from the PixeLINK camera

The second method to determine pixel size was used for the NybyOnline dataset. In this case the camera position is fixed for all images, imaging end faces at approximately the same position. Therefore, pixel size can be determined once and applied to all images. In this camera setup, the end faces are not imaged exactly from the front, but slightly from above, meaning that the end face is rotated compared to the image plane and perspective distortion occurs. Hence, the pixel size is not constant over the image. It is assumed in this model that each end face is positioned at the same place when imaged and that the end face is a flat and vertical surface, rotated around the horizontal axis compared to the image plane.

To compute the pixel size on the entire image a few parameters must be known: the distance between camera and object, the rotation of the object compared to the image plane, and the focal length of the camera lens. These parameters were estimated here according to (Tsai, 1986) using a few reference objects. The calibration method also includes parameters related to lens distortion, however these were not estimated here, since they are assumed to have a small influence on the imaging compared to other disturbances. The rough environment with a shaking conveyor together with end faces that are neither still nor perfectly vertical, are assumed to create larger errors than lens distortion.

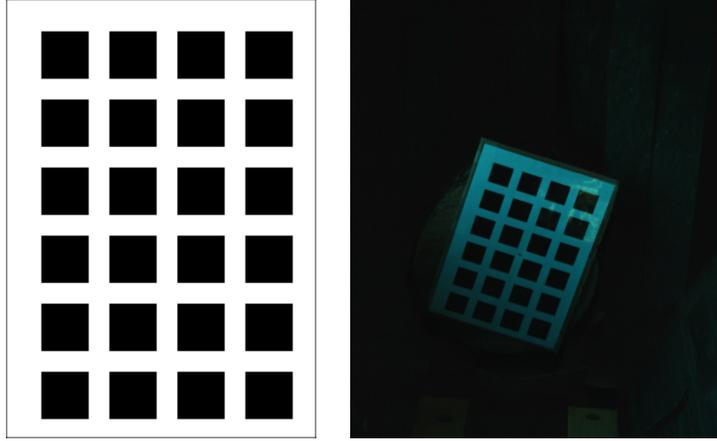


Figure 47: Reference $R_{(4,6)}$ in original (left) and imaged when nailed onto a log(right).

The calibration is performed by imaging a set of reference objects with known calibration points, which are then related to the corresponding points in the acquired images. Three different references were used; one large with 8×12 black squares denoted $R_{(8,12)}$, and two smaller with 6×8 and 4×6 squares, denoted $R_{(6,8)}$, and $R_{(4,6)}$, respectively. These were printed on paper, laminated, and glued onto cardboard. For the large reference the calibration points cover an area of 225×309 mm and for the smaller references 158×214 mm and 169×259 mm, respectively. The references were imaged while held still at the approximate imaging position for the end faces. Reference $R_{(4,6)}$ was also imaged as it was nailed onto a log and run on the conveyor. This reference and the acquired image is shown in Figure 47.

In camera calibration, image coordinates, $\mathbf{X} = (X_1, X_2)$ are related to real world coordinates, $\mathbf{x}^w = (x_1^w, x_2^w, x_3^w)$, via the camera coordinate system $\mathbf{x} = (x_1, x_2, x_3)$. A pinhole camera is assumed, with the coordinate system placed in the centre and the optical axis along \hat{x}_3 . The real world coordinates are transformed to camera coordinates through a rotation followed by a translation

$$\mathbf{x} = \mathbf{R}\mathbf{x}^w + \mathbf{T}, \quad (21)$$

where \mathbf{R} is a 3×3 rotation matrix and $\mathbf{T} = (T_1, T_2, T_3)$ is the translation along the different coordinate axes. The references are 2D objects with $x_3^w = 0$ and x_1^w and x_2^w given in mm with the origin in the reference centres. An illustration of the coordinate systems is shown in Figure 48.

The conversion from camera to image coordinates is determined by the focal length, f , and the distance between the pixel centres in the camera sensor, d ,

$$\begin{bmatrix} X_1 \\ X_2 \end{bmatrix} = \frac{f}{d \cdot x_3} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \quad (22)$$

here $d = 3.5 \cdot 10^{-3}$ mm (Anon. *PixeLINK*, 2004). The translation \mathbf{T} differs slightly between the imaged references, as does the rotation around all axes. Ideally, for the

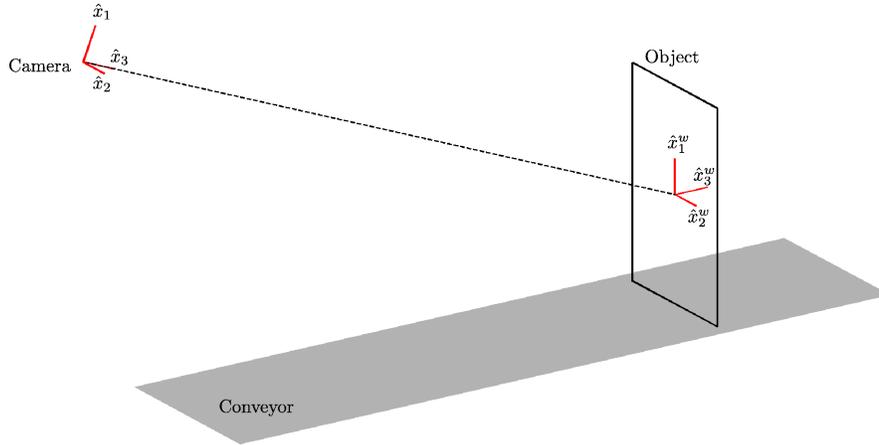
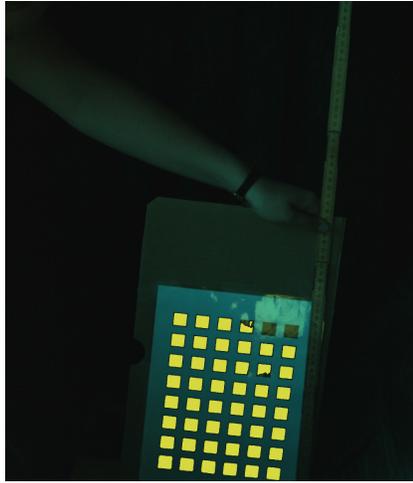


Figure 48: The camera coordinate system \mathbf{x} and the real world coordinate system \mathbf{x}^w in a schematic view of the NybyOnline camera setup.

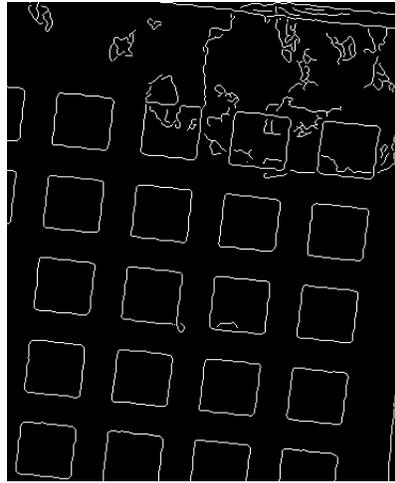
end face images there will be no rotation around \hat{x}_1 and \hat{x}_3 . The rotation around \hat{x}_2 and the translation T_3 will be constant, determined by the end face position in relation to the camera.

The calibration interest points are the corners of the black squares. These positions are known for the reference objects, but must be extracted from the images. Manual marking is demanding, tedious, and possibly error prone, therefore a semi-automatic method was developed and used. In this method the edges of the squares are used to determine approximate positions of the corners. Once the corner points are estimated, the result is refined by fitting straight lines to each row and column of the detected corners. The position of the square pattern is required in the extraction of interest points, therefore the user must identify the upper and lower, left and right squares initially.

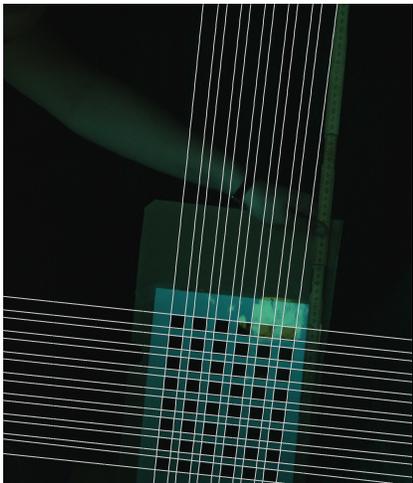
The corner points are estimated for one black square at a time, thus squares must be detected and labelled. The square detection was performed by smoothing the image with a Gaussian filter with standard deviation 2 to reduce noise and then perform global thresholding, with thresholds between 0.09 and 0.16 for the different images. After thresholding, all black regions were labelled. This procedure labels most of the black squares as well as the dark background and a few other regions. A manual removal of the erroneous labels was performed after which each label can be assumed to correspond to a square. The edges between white background and black squares are detected with Canny edge detection (Canny, 1986), using upper thresholds between 0.1 and 0.25, and lower thresholds as 40% of the upper ones. The parameters for edge detection and thresholding vary slightly between the different images. Suitable parameters were determined by visual inspection. The result after labelling and edge detection for the imaged reference $R_{(6,8)}$ are shown in Figure 49(a) and (b), respectively. The upper right part of the reference is reflecting incoming sunlight, causing difficulties both for labelling and edge detection. The edge detection is shown only for this part of the reference.



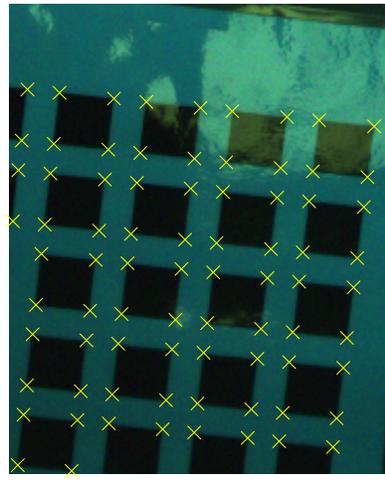
(a) Labelled image



(b) Detected edges for the upper right part of the reference.



(c) Straight lines fitted to the approximate corners.



(d) Final interest points (x) for the upper right part of the reference.

Figure 49: Different steps in the procedure of extracting interest points from reference $R_{(6,8)}$.

For each labelled square, four lines are fitted to the detected edge pixels. The approximate directions of the lines are known due to the known positions of the upper and lower, left and right squares, initially given by the user. The intersections between the lines are used as a first corner approximation. When all labels are visited the user must remove corner estimates that are clearly erroneous.

The performed corner estimation is sensitive to noise. To improve the result, straight lines are fitted to each row and column of detected corners in the image. The line intersections are the final detected interest points. This results in all corners being detected, which is not the case for the first approximate corner estimation. The fitted lines and the extracted interest points are shown in Figure 49(c) and (d), respectively.

The different reference objects were imaged at slightly different positions with different rotations compared to the image plane. However, the camera focal length is constant and is therefore estimated by least squares fitting of the results for the different reference images, resulting in $f = 18.0$ mm.

To compute the pixel size for all end face images, their approximate translation T_3 and rotation around \hat{x}_2 must be known. These are estimated from the reference that was imaged when nailed onto a log. This image is assumed to have the most similar position and rotation compared to real end faces. The estimated translation was $T_3 = 1965$ mm and rotation is -0.39 rad or -22.1 degrees. It is assumed that no other translation or rotation is present in the end face images. These parameters give a pixel width of 0.38 mm at the centre row of the image and 0.36 mm and 0.41 mm in the top and bottom, respectively. The pixel height is computed to 0.38 mm. This means that the pixels, although being quadratic in the image, does not cover quadratic regions on the end face.

To visualize the effect of this transformation, two synthetic patterns are shown in Figure 50 together with the result after applying the estimated distortions. The thin circles are superimposed on the circular patterns as a reference for easier comparison. The circular pattern is shown to illustrate the effect on annual rings. The largest visible effect occurs above and below the image centre, in the areas where the rings are almost horizontal. The perspective distortion compresses the object in height, meaning that when rings are narrower it will be difficult to resolve them.

The correct measurement area can now be computed by determining the real world coordinates for each pixel, using Equation 21 and 22. The Euclidean distance from each coordinate to the real world pith position is computed in mm and stored as a distance map. This information is then used to determine the interesting region for annual ring measurements, i.e., the region with distance 20 to 80 mm from the pith.

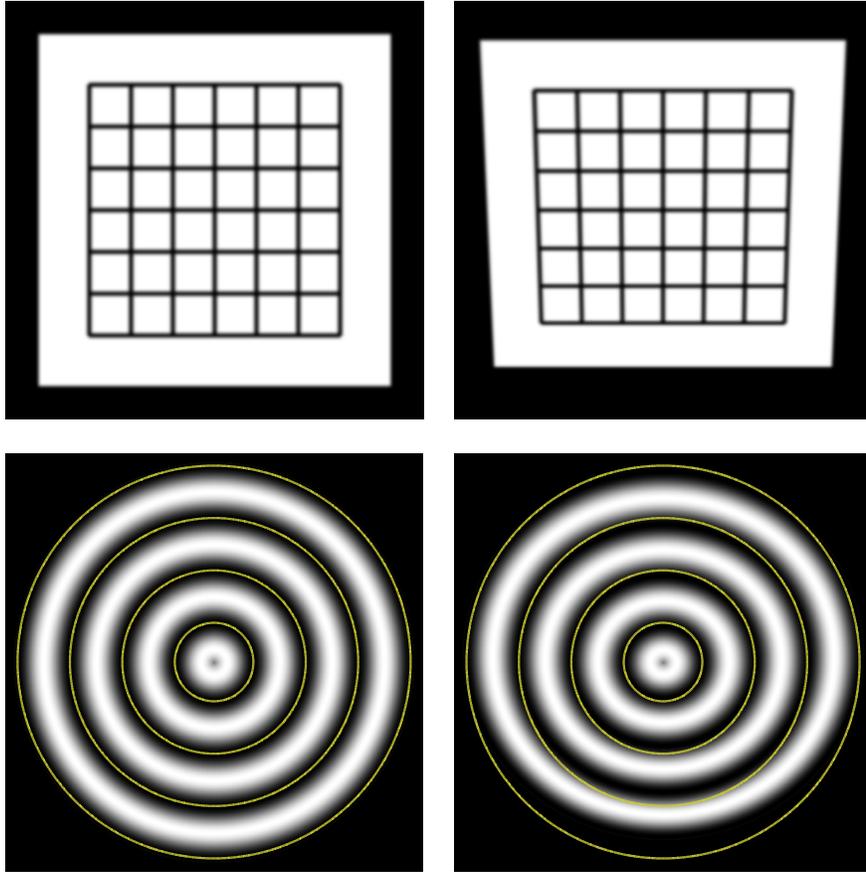


Figure 50: The effect of the estimated perspective distortion for the NybyOnline dataset, illustrated with two synthetic patterns (left) and the distorted versions (right). The thin circles are superimposed on the circular pattern for easier comparison of the images.

7.2 Automatic cropping

In the NybyOnline images a large amount of the content is background. The background is not interesting for the analysis of the end face features, thus reducing it, making the images smaller, will speed up computations. Due to similar properties in all images, it is possible to perform automatic cropping. The end face is completely included in a square region, the *bounding box* limited by the coordinates (x_1, y_1) , for the upper left corner and (x_2, y_2) , for the lower right corner. The aim is to automatically determine estimates (\hat{x}_1, \hat{y}_1) and (\hat{x}_2, \hat{y}_2) for the bounding box corners to perform cropping of the image, removing as much as possible of the background while keeping the entire end face.

The cropping is performed on an intensity image, I computed from a modified version of Equation 3, as discussed in Section 5.1. The end faces are rather

bright and the background is mainly dark, which implies that the end faces can be detected by thresholding the intensity. However, bright regions of the background exist, making this procedure insufficient. The bright regions originate mainly from three sources:

- A curtain of black plastic stripes is hanging behind the log. In some cases light is reflected in the curtain, creating a bright area in upper part of the image. This is mainly the case for small end faces, since larger ones occlude the curtain.
- A small part of the conveyor is visible below the end face. The conveyor belt consists of wooden planks lengthways and crossing transporters. The transporters are dark but the planks have similar colour as the end faces, resulting in one or two bright objects in the lower part of the image. The visibility of the planks in the images depends on if they are occluded by a transporter or not.
- The end faces are imaged slightly from above, as discussed in Section 3.2.6, meaning that parts of the log envelope surface are visible in the image. Depending on bark coverage and log brightness this can produce a bright area just above the end face in the image.

An example of a typical image I is shown in Figure 51. Note that both the curtain and the conveyor belt planks are visible as bright objects in the image.

The estimation of the corners (x_1, y_1) and (x_2, y_2) is performed in two steps: first for the vertical coordinate y and then for the horizontal coordinate x . The vertical cropping removes the curtain and conveyor, making the horizontal cropping straightforward. The possibly bright area created from the log envelope surface is not handled specifically here, but is assumed to influence the size of the bounding box only to a small extent. A lower limit for end face size is required for the automatic cropping to perform well. Here it is chosen as $\Lambda = 300$, corresponding to a bounding box of approximately 15×15 cm.

7.2.1 Vertical cropping

The vertical boundaries, y_1 and y_2 are estimated by examining the image row-wise. The mean value for each image row is determined,

$$\mu_x(j) = \frac{1}{M} \sum_{i=1}^M I(i, j), \quad (23)$$

where M is the number of columns in the image. A smoothed version $\bar{\mu}_x$ of the mean value is computed using a mean filter of length $2 \cdot w + 1$, with $w = 25$. To handle filtering at signal boundaries, the w first and last values of μ_x are mirrored in the first and last index, respectively, and the original part of the signal is filtered. Finally, the difference signal is computed approximating the derivative, i.e., $\Delta \bar{\mu}_x(j) = \bar{\mu}_x(j+1) - \bar{\mu}_x(j)$ for $j = 1, 2, \dots, N-1$, where N is the number of rows in the image.

The global maximum in $\Delta\bar{\mu}_x$ corresponds to the largest increase of intensity between two pixel rows. In most images the largest increase appears at the transition from background to end face, but in rare cases also the conveyor belt cause a large derivative. To detect the transition between background and end face the maximum in $\Delta\bar{\mu}_x$ from index 1 to index $N - 1 - \Lambda$ is detected, i.e.,

$$\nu = k \mid \Delta\bar{\mu}_x(k) = \max_j(\Delta\bar{\mu}_x(j)), \quad (24)$$

for $j = 1, 2, \dots, N - 1 - \Lambda$. Since the last $\Lambda = 300$ pixel rows are ignored, detecting the conveyor belt is avoided.

The fact that an end face is an approximately circular object means that the row-wise intensity change between background and object is not always sharp. The position ν therefore often appear at $y_1 + \delta$, where $\delta > 0$. This is the case in the image shown in Figure 51, where the dotted line marks the position of index ν . To approximate y_1 all local minima in $\bar{\mu}_x$ from index 1 to index ν are detected. The index for the local minimum closest to index ν is chosen as \hat{y}_1 .

The other vertical coordinate, y_2 is approximated as the index for the global minimum of the signal $\bar{\mu}_x$ from index $\hat{y}_1 + \Lambda$ to N . Indices \hat{y}_1 and \hat{y}_2 are shown in Figure 51 as solid lines. The ground truth in this image is $y_1 = 1196$ and $y_2 = 1595$, which were estimated to $\hat{y}_1 = 1161$ and $\hat{y}_2 = 1599$.

7.2.2 Horizontal cropping

The estimation of the horizontal boundaries, \hat{x}_1 and \hat{x}_2 are computed using the image I_{crop}^y , i.e. the image I cropped at the vertical indices \hat{y}_1 and \hat{y}_2 . Each column mean is computed and smoothed, similar to the vertical case,

$$\mu_y(i) = \frac{1}{N_{crop}} \sum_{j=1}^{N_{crop}} I_{crop}^y(i, j), \quad (25)$$

where N_{crop} is the number of rows in I_{crop}^y . A typical μ_y and the smoothed version $\bar{\mu}_y$ are shown in Figure 52 together with I_{crop}^y . It is clear that the indices corresponding to the position of the end face have considerably higher intensity than others. This is the case, since the curtain and planks are already cropped from the image, and thus have no influence on the end face detection. The detection could thus be performed using global thresholding. However, finding one threshold suitable for all images is difficult. Here, another approach is taken where an automatically determined threshold is used to detect the end face roughly, after which it is possible to make a more exact detection of the boundaries.

The threshold, T is chosen as the mean value of the global minimum and maximum in $\bar{\mu}_y$. Thresholding $\bar{\mu}_y$ at T gives a binary signal with ones corresponding to indices where the end face is present. Now, the task is to detect the boundary indices of the end face. Let the first and last indices corresponding to ones in the thresholded signal be τ_1 and τ_2 , respectively. Index \hat{x}_1 is determined as the last local minimum in $\bar{\mu}_y(i)$ for $i = 1, 2, \dots, \tau_1$, and similarly \hat{x}_2 is the first local minimum in signal $\bar{\mu}_y(i)$ for $i = \tau_2, \tau_2 + 1, \dots, M$. This way of determining the cropping indices

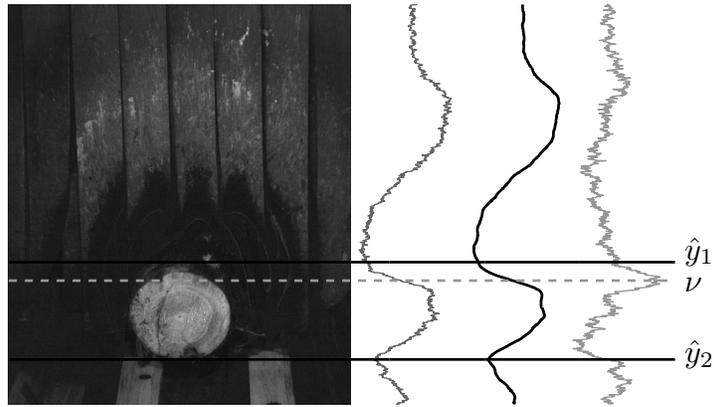


Figure 51: Original intensity image (left) together with signals μ_x (middle left), $\bar{\mu}_x$ (middle right) and $\Delta\bar{\mu}_x$ (right). The dotted line marks the index ν for the maximum value in $\Delta\bar{\mu}_x$. Solid lines show \hat{y}_1 and \hat{y}_2 .

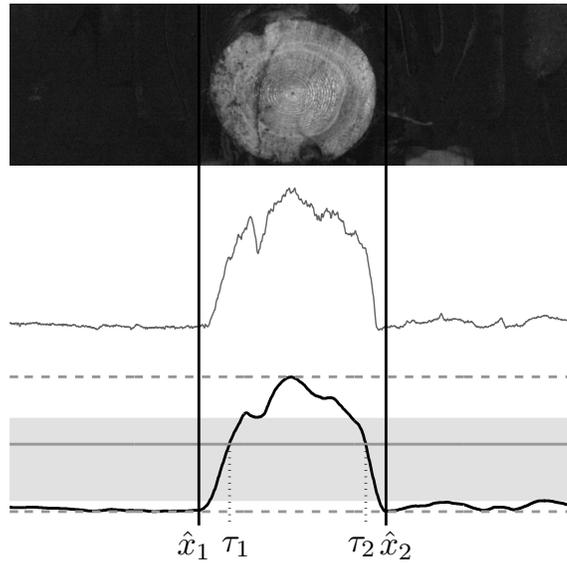


Figure 52: Image I_{crop}^y (upper) with signals μ_y (middle) and $\bar{\mu}_y$ (lower). The horizontal lines show the threshold T (solid) as well as the maximum and minimum values of $\bar{\mu}_y$. The vertical lines show \hat{x}_1 and \hat{x}_2 . The grey area marks the region for the values of T that would have given the same estimates for \hat{x}_1 and \hat{x}_2 .

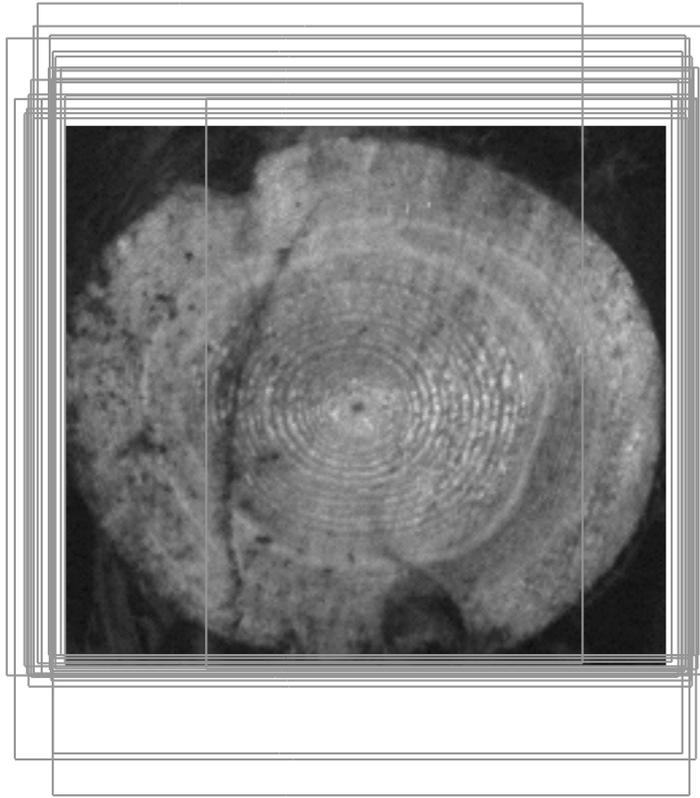


Figure 53: Illustration of the result for a set of 20 evaluation images. Each rectangle corresponds to the result of cropping one image. The shown image is 445×400 pixels.

is not sensitive to the threshold T . The grey area in Figure 52 illustrates the span of thresholds that would give the same cropping result. The estimated \hat{x}_1 and \hat{x}_2 are shown as vertical lines in Figure 52. The ground truth is $x_1 = 554$ and $x_2 = 998$, and the estimated indices were $\hat{x}_1 = 515$ and $\hat{x}_2 = 1024$.

7.2.3 Results

The performance of the suggested method is shown in Figure 53 for a small evaluation set of 20 images which is a subset of the larger evaluation set \mathcal{S}_E . The shown end face represents an arbitrary end face cropped manually at (x_1, y_1) and (x_2, y_2) . Each of the drawn rectangles illustrates the result for one of the images in the evaluation set normalized to the size of the shown image. That is, the shown boundaries are multiplied with the ratio between each image size and the size of the shown end face. In two of the illustrated results large parts of the end face is cut in the horizontal parts of the end face is cut in the horizontal cropping. These end faces were both very dirty with dark parts that were incorrectly detected as background in the automatic cropping. The ground truth was

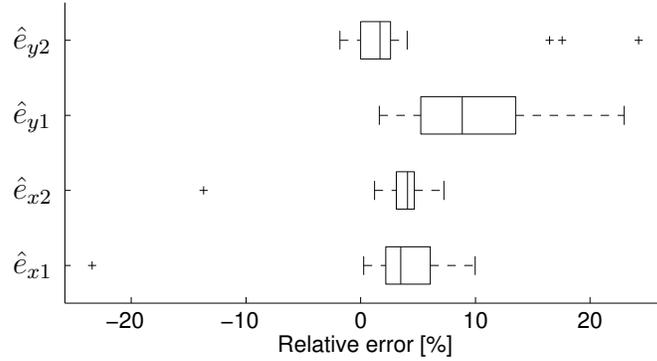


Figure 54: Results for the four different estimates for the evaluation set. See the text for explanation.

established manually by marking with a graphical tool the upper left and lower right corners of the region including the end face.

The results are also shown as box plots in Figure 54. Here, the relative error for each end face is computed as

$$\hat{e}_{y1} = \frac{\hat{y}_1 - y_1}{2(y_1 - y_c)}, \quad (26)$$

for \hat{y}_1 and similar for the other estimated coordinates, where y_c is the vertical centre of the bounding box, i.e., $y_c = (y_2 + y_1)/2$. This means that a positive value corresponds to cropping outside the end face and a negative value means that a part of the end face is removed. The lines in the boxes mark the median value, the lower and upper limit of the box is the 25 and 75 percentile, respectively and the whiskers cover all results that are not considered outliers (+). One large negative error is seen for both variables \hat{e}_{x1} and \hat{e}_{x2} , which are equal to the discussed large errors seen in Figure 53. The upper vertical coordinate \hat{y}_1 is the estimate with the highest median, and the largest errors, excluding outliers. This is an effect of the smooth intensity transition between background and end face, as well as bright parts of the log envelope surface being visible. The coordinate for the lower boundary of the end face \hat{y}_2 has three positive outliers, which means that a part of the conveyor is included in the cropped images.

When cropping the images it is more important that most of the end face remain in the cropped image, than that the cropping is as close to the true bounding box as possible. Therefore, positive errors in Figure 54 are desired rather than negative errors. The extra rows or columns of pixels that remain in the cropped image are not a problem.

7.3 Automatic removal of marks from uneven sawing

Marks from harvesting are clearly visible across the end face in many images. An example of an image with such a pattern is shown in Figure 55. The marks are disturbing both when it comes to estimation of the pith position as well as in the computations of the number of annual rings since they are detected in the computations of local orientations. An automatic method for removing them was presented in Paper VII and applied to the images before the further analysis. By automatically detecting the direction of the marks, a suitable filter mask is created in the Fourier domain and applied to the image spectrum.

The repetitive pattern of lines or marks that is the effect after uneven sawing is often visible in the Fourier domain as a region with high energy. By detecting the region, it is possible to construct a filter that suppress the energy, and thus removes the marks. A small part of the Fourier spectrum for the image in Figure 55 is shown in Figure 56. This part is the centre of the Fourier spectrum, which is the most interesting part when it comes to these marks. The energy in the centre was set to zero and the spectrum was smoothed with a Gaussian filter. The pattern from uneven sawing creates higher energy in the Fourier spectrum close to the horizontal axis. By measuring the intensities in different directions from the centre, the region with high intensity can be detected, as shown in Figure 56. The dashed vertical line is the position for the detected maximum in the signal, representing the direction of the marks.

A filter is created as a line in the Fourier domain, interrupted in the spectrum centre and smoothed with a Gaussian filter. The interruption avoids removal of the lowest frequencies in the image, which is not required here. The filter is shown in Figure 57 together with the result after filtering, i.e., pointwise multiplication, the spectrum in Figure 56. Note that the shown result is only an illustration with the centre energies removed and the smoothed spectrum for easier comparison with Figure 56. In the reduction of the marks, the original Fourier spectrum is the one filtered. The result after filtering the image in Figure 55 is shown in Figure 58.

The filtering is performed as a preprocessing step for all images in the Heby2006, Heby2007 and NybyOnline datasets prior to further analysis. Some of the end faces do not need this filtering, but no automatic method was developed to determine which one required it, thus it was applied to all. The areas with annual rings parallel to the filtered pattern is slightly affected by the procedure. However, it was concluded that performing the filtering was important enough to take the risk of affecting the annual rings.



Figure 55: Original image with marks from uneven sawing.

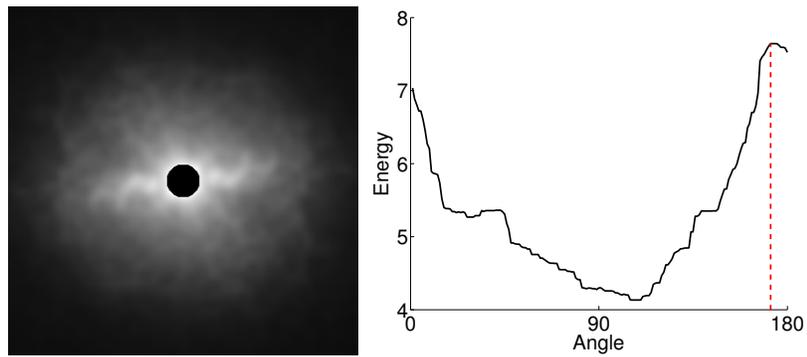


Figure 56: The centre part of the filtered Fourier spectrum (left) and the energies corresponding to different angles from the centre (right). The angle is measured clockwise from the horizontal axis.

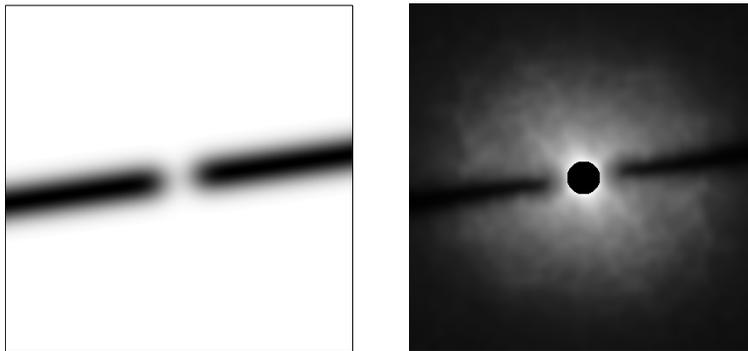


Figure 57: The constructed filter (left) and the result after filtering (right).

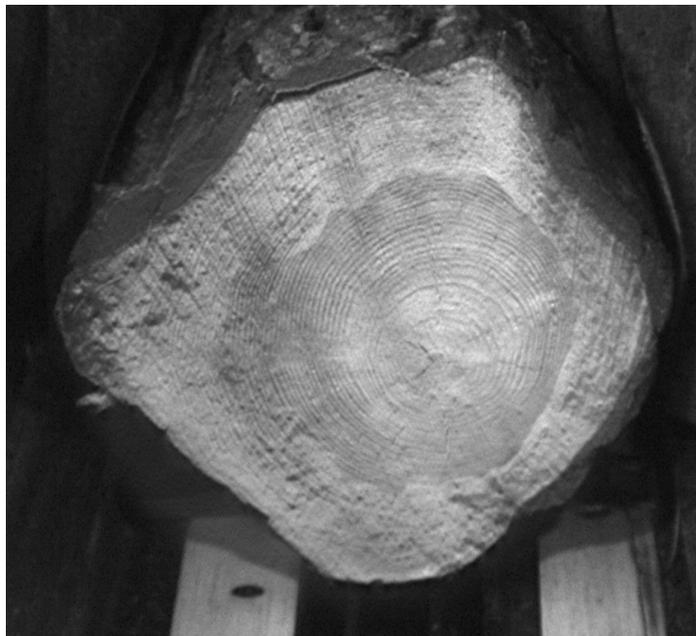


Figure 58: Result after filtering the image in Figure 55.

8 Analysis of end face features

So far, the presented work includes background, image acquisition, theoretical tools and preprocessing of the images, which are all needed for the last step: the actual analysis of the end face features described in this chapter. The extracted features are the pith position and the number of annual rings. The method for pith detection was presented in Paper V, and is used as a prerequisite for the annual ring measurements. The annual ring computations were described and evaluated in Paper VI and Paper VII for pine end faces. New results for pith detection are presented here for images that were not analyzed before. Also, initial results for computing the number of annual rings in spruce are presented. The methods are briefly described here. Details can be found in the corresponding papers.

8.1 Pith detection

The position of the pith in a log is related to wood features. In the case of compression wood in the stem the pith is often eccentric (Bowyer et al., 2007). Even if no compression wood exists an eccentric pith means that the tree growth has been uneven in different parts of the log, which is undesirable in wood products. If the shape of the end face is also outlined, the position of the pith can be compared to the centre of mass of the end face, as a measure of eccentricity. This measure was computed in Paper V, but the aim of estimating the pith position in this work has been to use the result as a starting position for the annual ring measurements. This is common to many methods considering annual rings on end faces or cross sections of wood logs.

The approach for pith detection is similar for many methods, independent of the imaging device or quality of the end face surfaces for the considered images. The goal is to detect the centre of the annual rings, using the local orientations of the rings in the computations. If the rings were completely circular and concentric, the local orientation of each pixel would be directed towards the ring centre, i.e., the pith. Therefore, by computing two such orientations the point where the lines intersect would be the position for the pith. However, a few difficulties occur with this model. The first is that the annual rings are neither perfectly circular nor completely concentric. However, it is probable that most of the orientations will point rather close to the pith. Thus, by combining information from many orientations it is possible to make an estimation of the pith position. The second difficulty is that the local orientations of the annual rings can be hard to estimate, due to rough end faces and noisy images. The method for orientation estimation must therefore be robust enough for the considered end faces and images.

The described approach is used in several papers, e.g., Andreu and Rinnhofer (2001); Bhandarkar et al. (1996); Entacher et al. (2007); Hanning et al. (2003); Longuetaud et al. (2004), and Österberg et al. (2004). Other papers suggest different methods, but with certain parts still similar. Although the approach is the same, the methods are not equal. Computations of the orientations differ, and so does the method for combining them into an estimation of the pith position.

Computations of the local orientations have been performed by thresholding (Bhandarkar et al., 1996; Entacher et al., 2007; Laggoune et al., 2005), gradients (Longuetaud et al., 2004; He et al., 2008), and edge detection (Chalifour et al., 2001). In these methods, parts of the annual rings are detected, and thus the local orientation. However, they are all applied to rather high quality end faces and images acquired either using CT scanner or through a microscope or by a desktop scanner.

Another approach is to take into account a larger area when determining the local orientation, which makes it more robust to rough surfaces and noisy images. This can be performed using the fast Fourier transform (FFT), which is the case in Hanning et al. (2003) and Österberg (2009) for planks and end faces, respectively; or Gabor filters, which was used in Andreu and Rinnhofer (2001). An alternative is to use separable filters and linear symmetry to determine local orientation (Bigün, 1992), which was used in Paper V. There, orientations are determined in different frequency scales using a Laplacian pyramid. The local orientations are computed using Gaussian filters, which is fast due to the possibility to separate the filters in spatial coordinates. In Figure 59 a small part of an image is shown, where the local orientations were computed both using the gradient and the Laplacian pyramid. The two first levels are shown from the pyramid, which are the levels that were used in this work. The end face is too rough for the gradient to perform well and the first level of the pyramid is also sensitive to the rough surface. The second level of the pyramid, however, performs very well in the shown image.

The difference in performance between thresholding and detecting local orientation using FFT is discussed in Hanning et al. (2003). There, the pith is detected on the front side of lumber. With polished front sides, thresholding outlines parts of the annual rings which are then used to estimate the pith position. For unpolished sides this does not work, thus the FFT is used instead. Österberg (2009) concluded that global thresholding was not useful in detecting the rings on sawn end faces. These conclusions fit well with what was discovered in this work.

In the methods for computing the pith position, local orientations are often combined to a voting matrix or Hough accumulator image. In Österberg (2009) two pixels and their corresponding orientations are picked from the end face, and the local orientations are computed. An estimate for the pith position is taken as the intersection point between the orientations. Two new points are then selected and a new position is estimated. This is iterated until the difference between two estimates is small enough. In other methods several orientations are combined into a voting matrix e.g., Andreu and Rinnhofer (2001); Entacher et al. (2007); Hanning et al. (2003), and Longuetaud et al. (2004), which was also the approach in Paper V.

It should be mentioned that there exist other approaches for pith detection, e.g., in Flood (2003) where whorls of knots are used to detect the pith position in CT images, or Sjöberg et al. (2001) who determines the pith position on planks by comparing local orientations to pre-computed directions. A comparison between different methods such as combing the gradient directions, measuring curvature and using techniques for determining minutiae in fingerprint images was presented in Entacher et al. (2008) for pith detection and it was concluded that taking the intersections between gradients gave the best results.

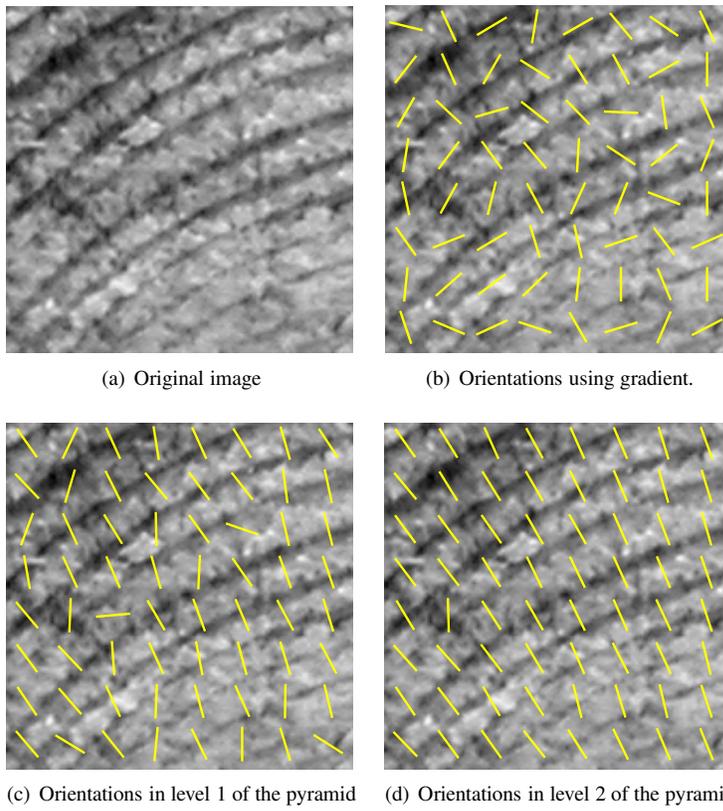


Figure 59: Detection of the local orientations of the annual rings computed using gradients and two levels of the orientation pyramid.

Two versions of pith detection were described in Paper V, differing mostly in the method for local orientation estimation, using quadrature filters (Granlund and Knutsson, 1995) or a Laplacian pyramid. Both performed well for the analyzed images, i.e., the Forssjö, Ultuna, and Nyby2005 datasets. The second version, using the Laplacian pyramid, has been used since then as it computes local orientations faster than the quadrature filters. The method is robust for rough surfaces, but with this type of data, i.e., untreated end faces, cases occur when the method performs bad or fails. Therefore, a validity measure is computed for the estimated pith, which can give a warning in case of an uncertain pith estimate.

8.1.1 Additional experiments

The pith position has been computed here for end faces in datasets that were not analyzed in Paper V. This includes both the Heby2006 and Heby2007 datasets and the evaluation set, \mathcal{S}_E from the NybyOnline dataset. The pith detection results for the dataset used in Paper VII are also shown here. In Paper VII the pith position was used as a prerequisite for annual ring measurements, but the performance of the pith estimation was not examined separately. In Paper V, intensity images computed using Equation 4 were used. However, when applying the pith detection to the NybyOnline data, another conversion is used (equal to the one in the automatic cropping described in Section 7.2), since it was more suitable for annual ring measurements.

The ground truth for pith position was established by the author by choosing one pixel in each image using a graphical tool. The pith is considered to be one pixel, although the physical pith on an end face is larger. In a few end faces the position of the pith is very difficult to determine, e.g., if the area around the pith is covered by dirt or in case of severe rot. However, ground truth was still established as well as possible.

The method used for automatic pith detection consists mainly of two steps: computing local orientations using Laplacian pyramids and combining them into an estimation of the pith position in an intersection step. The orientation estimation requires three parameters, which are listed in Table 6 for the considered datasets. The parameter σ_1 is the standard deviation for the filter creating the Laplacian pyramid, i.e., it determines the frequency of the signals in each level of the pyramid. Here σ_1 is chosen so that the annual ring pattern is caught in the two first levels of the pyramid. For higher levels the frequencies are assumed to be too low. The parameter β determines the size of the region that is considered in the estimation of local orientation. A large value takes a larger neighbourhood into account, which is robust to noise. However, the larger the neighbourhood, the more does the annual ring pattern deviate from approximating a simple signal, and the estimated orientation can therefore be uncertain. The final standard deviation, σ_2 is the size of a Gaussian filter, smoothing the detected local orientations. In the computations of local orientations using this method, a certainty value is given for each computed orientation. This certainty value is used in the intersection step of the pith estimation.

The values of the parameters listed in Table 6 were chosen by visual inspection of the first two levels of the Laplacian and orientation pyramids of a few images in

Table 6: Parameters for computing the local orientation.

Dataset	σ_1	β	σ_2
NybyOnline	2.0	1.5	1.5
Heby2006	3.0	3.0	3.0
Heby2007	3.0	3.0	3.0

the \mathcal{S}_E set and Heby2007 set. The parameters chosen for \mathcal{S}_E were also used for the data in Paper VII, which are here listed as NybyOnline, since both datasets are subsets of the NybyOnline set. For the images in Heby2006, it was assumed that the same parameters could be used as for Heby2007.

In the intersection steps of the method, the orientations are combined into an estimation of the pith position. The number of local orientations that are taken into account is determined by two parameters: a threshold T_c and the size of a moving window with $B \times B$ pixels. In each non-overlapping window, the orientation with the largest certainty value is chosen as a representative for the region, if the certainty is above the threshold T_c . A line with the same direction as the local orientation is drawn in an intersection image, representing possible pith positions. For each new examined region, the lines are accumulated in the intersection image. A Gaussian filter with standard deviation σ_I smooths the intersection image, after which the weighted mean of the largest values is taken as an estimate of the pith position. A threshold, T_S sets the limit for which values are taken into account in the mean computations. The intersection step is performed three times: once for the complete orientation image in level 2 of the pyramid, and afterwards for a small area around the first result, both in level 1 and level 2. This means that in total three estimates for the pith position is computed, where the first is assumed to be a rough estimate since it takes into account local orientations from the complete image.

The parameters used are listed in Table 7, and are equal for all datasets. By choosing $T_c = 0$, there are no restrictions on the certainty values of the orientations taken into account. The size of the moving window B is chosen so that it is rather large for the first intersection step, when the complete image is taken into account, and smaller for the second intersection. In the second step the size of the moving window doubles when the resolution doubles (from level 2 to level 1 in the pyramid.) The same goes for the standard deviation σ_I . With threshold $T_S = 1 - \varepsilon$, where $\varepsilon = 10^{-10}$, in practice no mean is computed of the largest values in the intersection image, but the pith position is chosen as the position for the largest value. For the intersection step, the parameters were determined using 20 images from the Heby2007 dataset, which is 10% of the complete dataset. The error for each step was compared to the ground truth after different steps. The chosen parameters provided a result with a low median error and among the smallest variations. The images used for training are included in the result.

The pith estimation result is chosen from one of the two estimates in the second intersection. The level (1 or 2) with the smallest standard deviation of the intersection image is assumed to have a narrow peak and therefore be a more exact estimate. The standard deviation is computed as a mean of the standard deviations for the hor-

Table 7: Parameters for the intersection steps.

	First intersection				Second intersection			
	T_c	B	σ_I	T_S	T_c	B	σ_I	T_S
Level 1	-	-	-	-	0	4	4	$1 - \varepsilon$
Level 2	0	8	2	$1 - \varepsilon$	0	2	2	$1 - \varepsilon$

horizontal and vertical directions, respectively, normalized with the size of the image. The standard deviation for the horizontal direction is computed as

$$\sigma_x = \sqrt{\frac{1}{M} \sum_{i=1}^M (x_i - \bar{x})^2}, \quad (27)$$

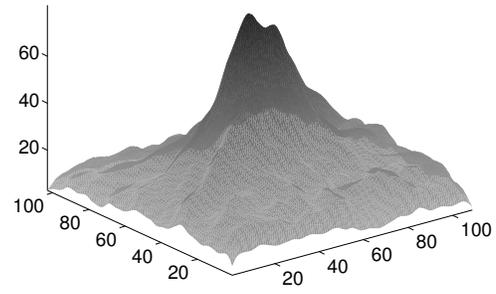
where \bar{x} is the mean value for the horizontal coordinate and M is the number of columns in the intersection image. The standard deviation σ_y for the vertical coordinate is computed in a similar way. The total standard deviation is

$$\sigma = \frac{1}{2} \left(\frac{\sigma_x}{M} + \frac{\sigma_y}{N} \right), \quad (28)$$

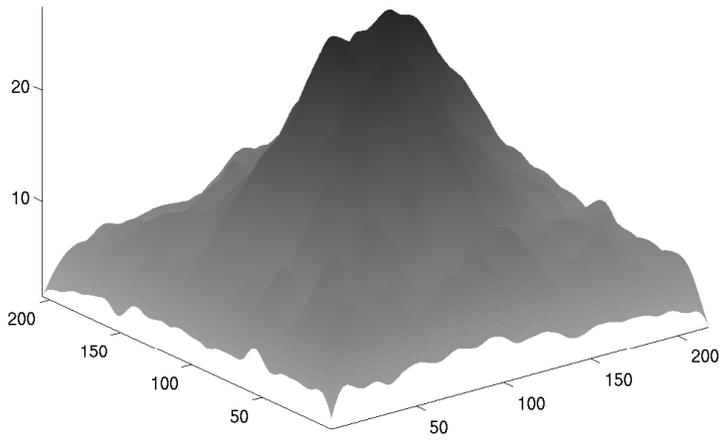
where N is the number of rows in the intersection image. Intersection images for level 1 and 2 for an image are shown in Figure 60. The standard deviation is smaller for the second level, thus the position for the maximum there is the final estimate of the pith.

In Paper V the computed standard deviations were used as validity measures for the detected pith. However, the interpretation of this measure is not straightforward, and it can be difficult to set a threshold for which results are valid or not. In Paper VII the validity is measured as the distance between the three different estimates of the pith position that are computed in the intersection steps. A large measure thus means long distances between the estimates, i.e., the estimate is deemed invalid. For a valid result, the different estimates will be similar, and thus the validity measure is low. The validity is measured in mm.

Results for pith estimates for S_E are shown in Figure 61, sorted according to the validity measure. A few end faces have large errors. The end face with the largest error was completely covered with mud, and the second largest was covered with snow to approximately 2/3 of the end face. All results with large errors also have a high value for the validity measure. The result is also plotted as a histogram in Figure 62, with the three largest errors removed. The results are shown in the end face images in Figure 63 for a few cases, sorted according to the validity measure. For the images on the top row the annual ring pattern is clear and no particular disturbances are present. The other end faces show a few difficulties, such as cracks in (c) and (e). However, in (c) the ring pattern is clear, thus the pith estimation is not disturbed by the cracks. In (e), knots are also present close to the pith, which results in a high validity measure, since both the cracks and the knots disturb the detection of local orientations. For the end faces in (d) and (f), the ring pattern is less clear due to a rough surface (d) and an end face slightly out of focus (f). In (f) the result is



(a) Level 2, $\sigma = 0.14$



(b) Level 1, $\sigma = 0.17$

Figure 60: Second intersection images for the two considered levels in the pyramid.

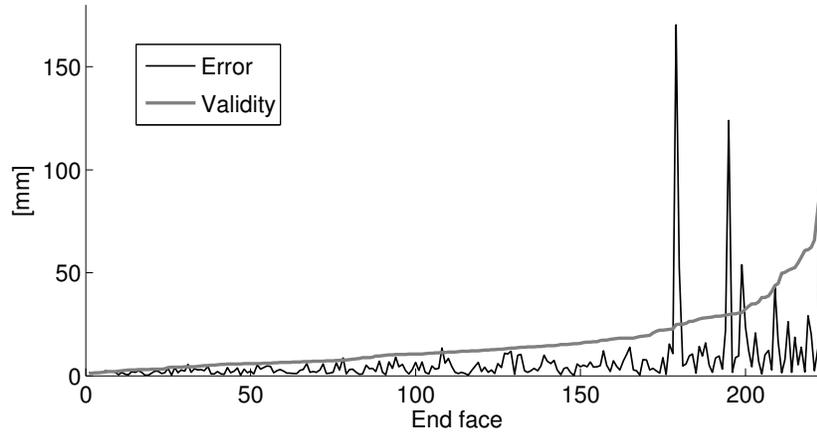


Figure 61: Result for the images in the evaluation set, \mathcal{S}_E sorted according to the validity measure.

still very good, although the validity measure is high. The result in (d) is the worst for the shown end faces.

The dataset that was collected for the evaluation in Paper VII consists of 70 end faces imaged three times: twice as they appeared at the sawmill, and once with the end faces newly sawn. A comparison for the pith detection method was performed using this data, with one set of images from the stored end faces compared to the newly sawn. The results are plotted in Figure 64, showing that the pith detection performs slightly better for the newly sawn end faces. This is an expected result, since they have less dirt, higher contrast, and smoother surfaces. However, the difference in performance is small.

Results for the Heby2006 set are shown in Figure 65 both as a histogram and plotted together with the validity measure. It can be seen that the large errors also have large validity measures. The results for the Heby2007 dataset are shown in Figure 66. A few results from both datasets are plotted in the respective images in Figure 67. Note that among the three images shown for the Heby2006 set, two are the ones that gave the largest errors for the complete set. Comparing the images with the ones shown in Figure 63 the images from the Heby sets have lower contrast and less intensity difference between earlywood and latewood. For Figure 67 (a) and (d) the ring pattern is clear and the results are good. Also the end face in (e) is handled well, although the marks from uneven sawing are disturbing. In (b) and (f) the ring pattern is not clear enough to determine the pith position, and in (c) the ring pattern is only visible for small parts of the shown image, which is also problematic.

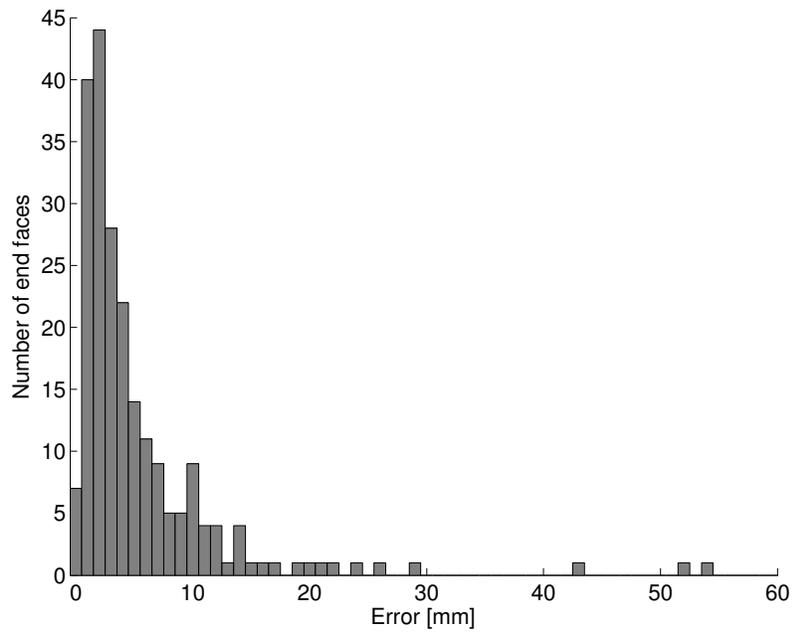
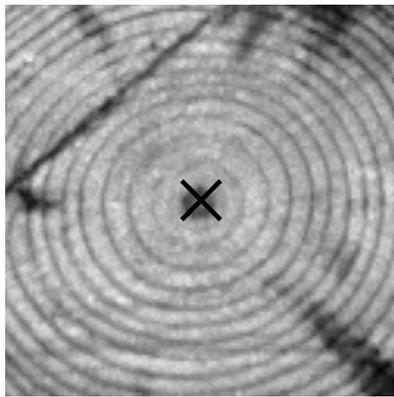
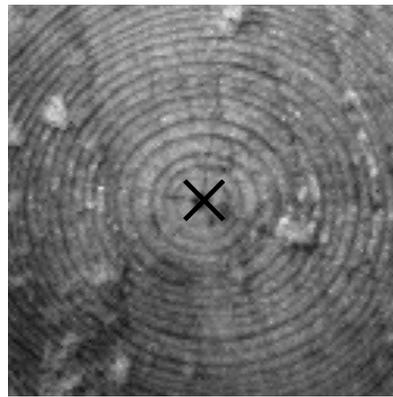


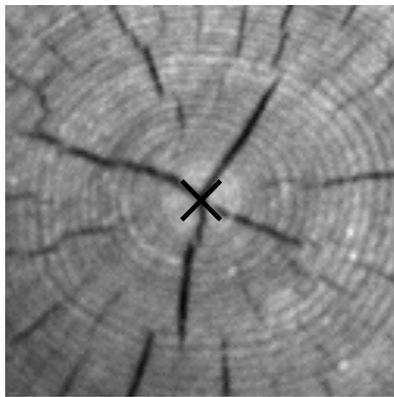
Figure 62: Results for the images in the evaluation set, \mathcal{S}_E . Three outliers have been removed.



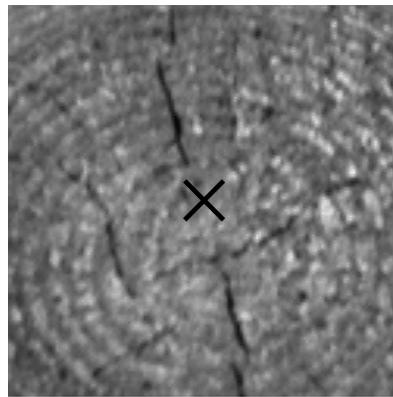
(a) $e = 1.4, v = 1.2$



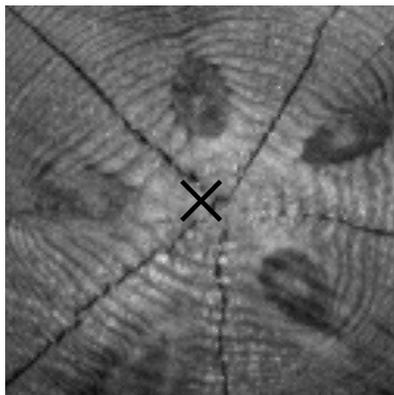
(b) $e = 1.6, v = 4.5$



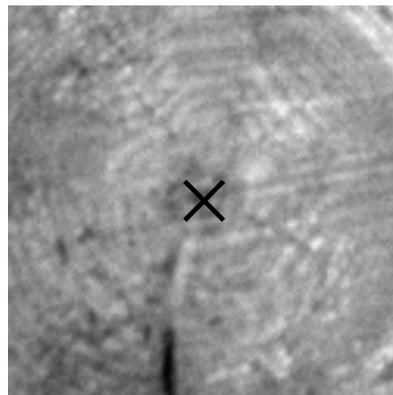
(c) $e = 0.4, v = 12.0$



(d) $e = 15.5, v = 22.5$

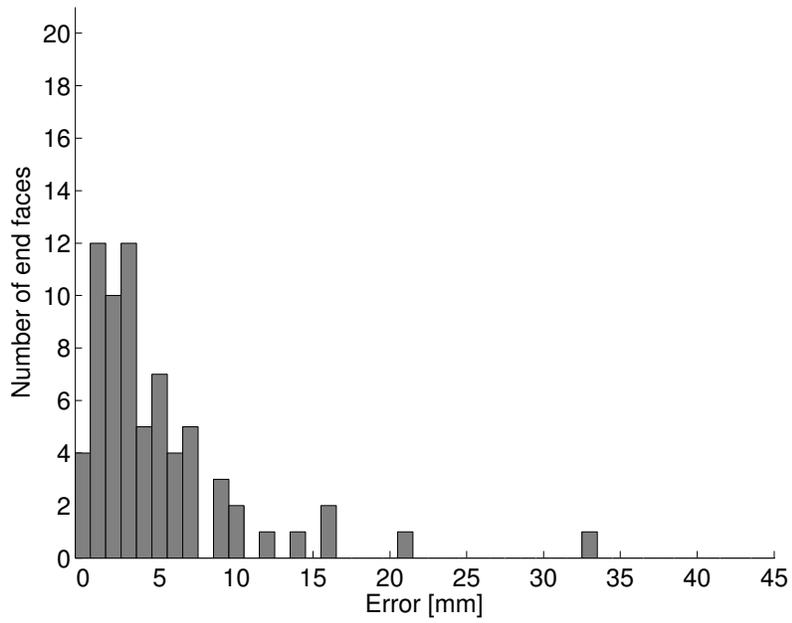


(e) $e = 5.7, v = 25.5$

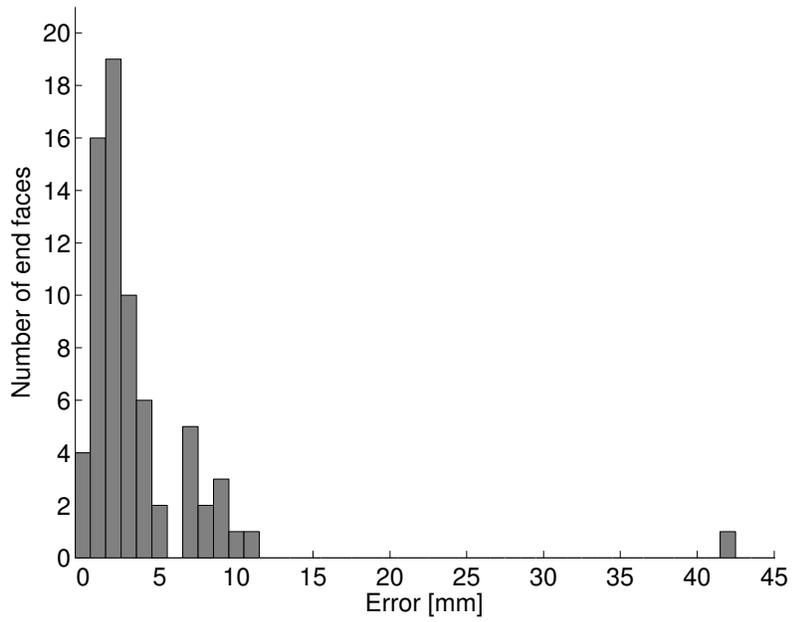


(f) $e = 1.6, v = 28.5$

Figure 63: Result for a few end faces in \mathcal{S}_E . The images are sorted according to the validity measure, v . The error e and validity measure are given in mm.



(a) Stored end faces



(b) Newly sawn end faces

Figure 64: Result for the pith detection for one set of end faces that was first imaged after storage, and then when newly sawn.

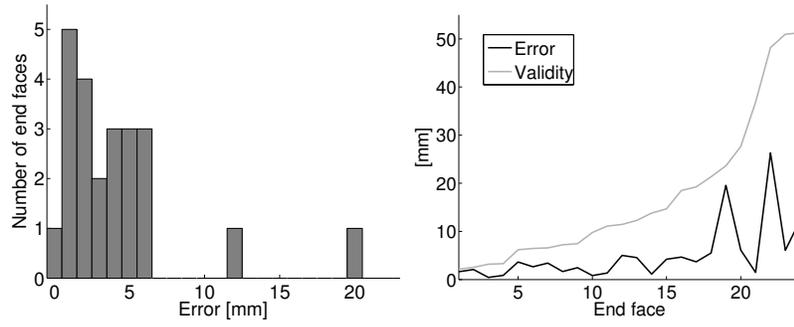


Figure 65: Result for the images in the Heby2006 set.

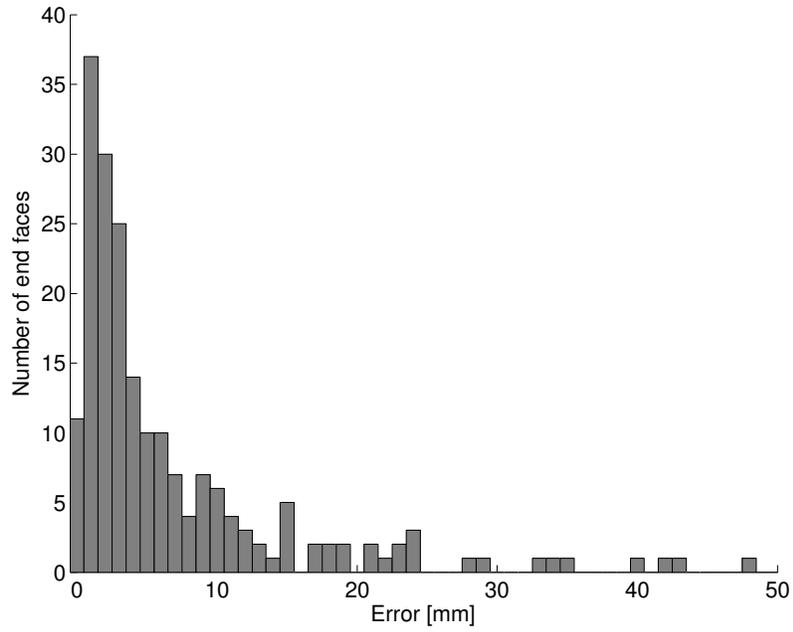
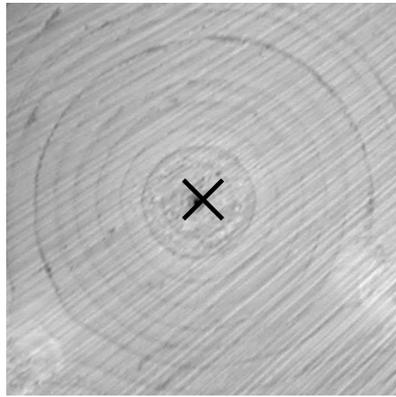
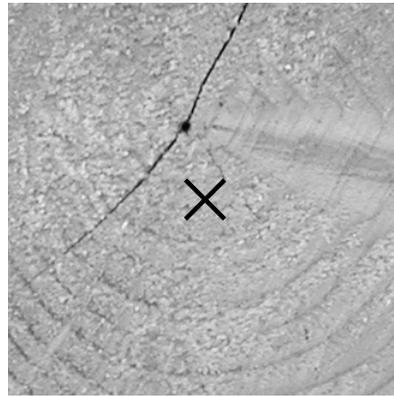


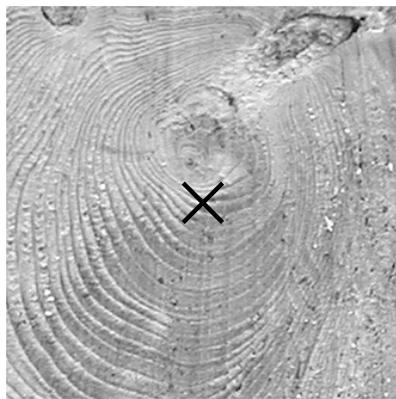
Figure 66: Result for the images in the Heby2007 set.



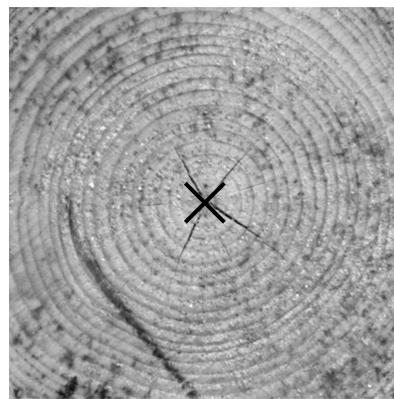
(a) $e = 1.6, v = 2.2$



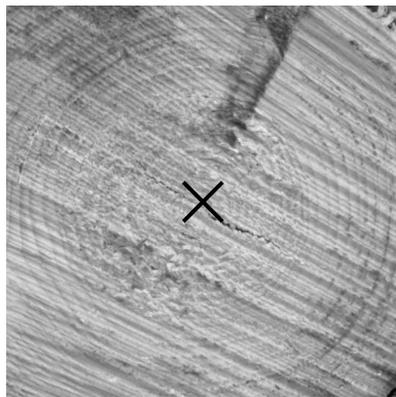
(b) $e = 19.5, v = 23.6$



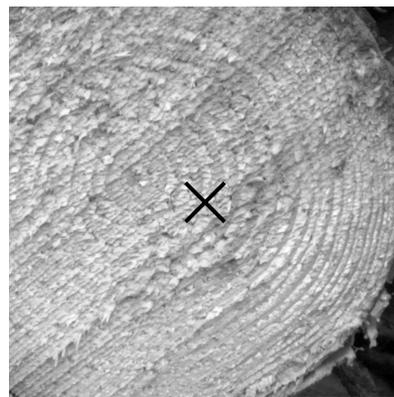
(c) $e = 12.3, v = 51.2$



(d) $e = 0.4, v = 1.4$



(e) $e = 1.4, v = 8.3$



(f) $e = 19.0, v = 35.2$

Figure 67: Result for a few end faces in the Heby2006 (a-c) and Heby2007 (d-f) sets. The images are sorted according to the validity measure, v for each of the datasets. The error e and validity measure are given in mm.

8.2 Annual ring detection

For annual ring measurements on end faces and cross sections, two application areas are the most common: quality measurements and dendrochronology. Different types of end faces and images give different possibilities for the measurements. Applications in dendrochronology, computing annual ring width requires determination of exact boundaries between different annual rings, and even occasionally between earlywood and latewood. To be able to measure such details, the quality of end faces and images must be very good, e.g., as in Soille and Misson (2001) where end faces are prepared and then imaged in a scanner to make it possible to determine annual ring area.

Annual ring contours can be outlined on high quality images such as MRI or CT data (Morales et al., 2004; Entacher et al., 2007) or scanned end faces (Laggoune et al., 2005) using methods based on edge detection and gradients. Lower quality data can also be useful if global information is taken into account as in Cerda et al. (2007) where the knowledge that two adjacent rings have approximately the same shape, is used. The method depends on the outer shape of the end faces being correctly outlined.

For end faces with a rough surface, or lower quality CT data, annual ring width has been determined using local frequency (Tsui, 1996; Andreu and Rinnhofer, 2001). In Hanning et al. (2003) and Österberg (2009), density maps of the annual ring width were computed on planks and end faces, respectively, using the FFT. It can, however, be difficult to determine a proper size for the window in which the FFT is computed. Tsui (1996) discusses the problem with few rings in a small block, and thus low angular resolution in the Fourier spectrum. A larger block means that the ring pattern deviates from a simple signal, and the angle and frequency is thus more difficult to determine.

Here, the focus has been on measuring the annual rings along a few rays from the pith and outwards. Ideally, the complete end face should be examined, but it was concluded that the end face quality is rarely good enough for this, and also that time can be saved by computing the annual rings only for a few directions, which should be considered for online computations in the sawmill.

To measure along one or a few rays has been considered earlier (Conner et al., 1998; Chalifour et al., 2001; He et al., 2008), but the methods applied there are not robust enough for untreated end faces. An attempt to measure the number of annual rings for quality analysis on untreated pine and spruce end faces was made by Jonsson (1992), where images of size 512×512 are acquired using a video camera. The annual rings were counted using the intensity along four different rays, and it was concluded that this method gave problems both for narrow annual rings, due to low resolution, and also for wide rings, where dirt or noise was often counted as annual rings.

In Paper VI a method to count the number of annual rings on end faces in sawmill industry was developed. This method was later modified and thoroughly evaluated in Paper VII. The grey weighted polar distance transform was used to compute the number of rings. The approach is the one shown in Section 6.4 in Fig-

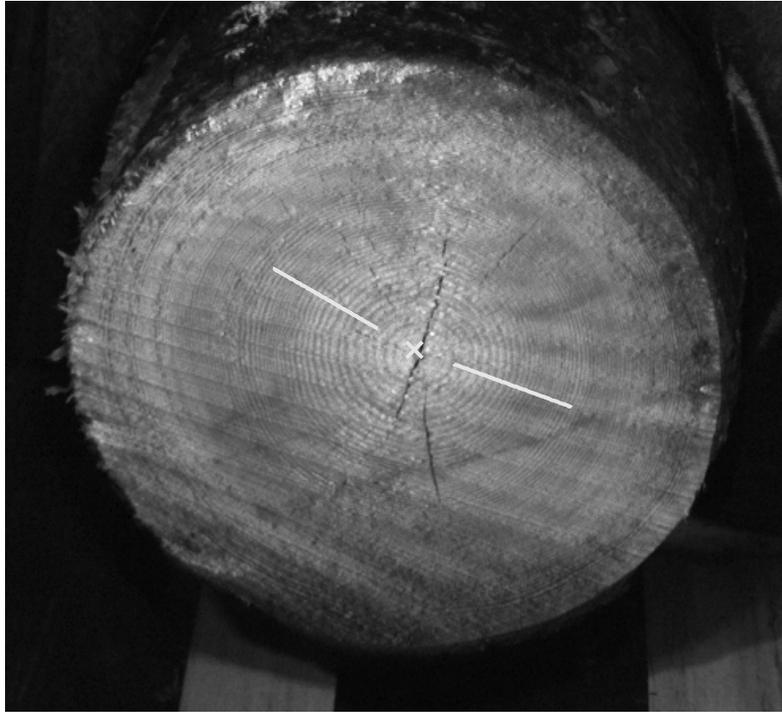


Figure 68: Automatically detected orientations for measuring the number of annual rings along two rays from 20 to 80 mm from the automatically detected pith (×).

ure 43 and 44, where the GWPDT is applied from a line, or ray, going from the pith and outwards, and the distance map is analyzed along another ray. The number of local minima along the analyzed ray in the distance map is ideally equal to the number of annual rings.

The direction for measurements is chosen automatically using the orientation pyramid computed in the pith detection step. The computed orientations are compared to the direction from the detected pith. An area where these are similar is assumed to have a ring pattern that is clear and is therefore chosen for measurements. No analysis of the performance of this approach was made, but it was concluded by visual inspection that the choice of measurement direction was satisfactory in almost all cases in the sense that the chosen direction was one with clearly visible rings (compared to other directions). By choosing two directions for measuring, as in Paper VII, the chance of picking not only an area with clear pattern, but also the correct measurement area is increased in comparison to measuring in only one direction. In Figure 68 the automatically detected directions are shown for an end face from the NybyOnline set.

The method for counting annual rings was evaluated on a dataset of 70 end faces. Each end face was imaged three times in the NybyOnline setup: twice as they appeared at the sawmill and once after a thin disc was sawn off to produce a

Table 8: Result for image series 1.

	True positives			False positives			False negatives		
	M1	M2	M3	M1	M2	M3	M1	M2	M3
A	56	55	47	3	4	4	2	3	11
B	8	6	7	2	3	11	3	5	4
C	0	1	1	1	1	0	1	0	0

Table 9: Result for image series 2.

	True positives			False positives			False negatives		
	M1	M2	M3	M1	M2	M3	M1	M2	M3
A	55	52	50	2	4	4	3	6	8
B	8	6	6	3	6	9	3	5	5
C	0	1	0	2	1	1	1	0	1

Table 10: Result for image series 3, newly sawn end faces.

	True positives			False positives			False negatives		
	M1	M2	M3	M1	M2	M3	M1	M2	M3
A	57	49	46	4	5	5	1	9	12
B	7	6	5	2	8	12	4	5	6
C	0	1	1	0	1	1	1	0	0

fresh surface. As the end faces were imaged, they were also graded by a grader and a grading inspector at the sawmill measurement station. Ground truth was also created for each end face by the grading inspector thoroughly counting the number of rings in the correct measurement area. The dataset consisted of 58 end faces of grade *A*, eleven end faces of grade *B* and only one of the worst quality class, grade *C*. The results for the measurements are listed in Table 8 to 10, where image series 1 and 2 correspond to the untreated end faces and image series 3 are the newly sawn. The proposed method is *M1* and the results for the grader and grading inspector are *M2* and *M3*, respectively. If all end faces are measured correctly, the true positives should be 58, 11 and 1, for grades *A*, *B*, and *C*, respectively. For the false positives and false negatives the ideal result is 0.

The results show that the automatic method performs as well as the manual estimation when classifying the logs into the quality classes as described in Section 2.4. The single end face with less than 12 rings (grade *C*) has not been correctly classified in any of the image series. Evaluation with more end faces in grade *C* must be performed to establish if this is a coincidence or a systematic error. Possibly, the method must be tuned slightly to measure these end faces correctly. It is noticeable that there is no improvement of the results for the end faces that were newly sawn, neither for the proposed method, nor the manual graders. It was expected that the automatic method could perform better for these images, however, since the method was developed using untreated end faces this result is not completely surprising. It is possible that the rather low image quality reduces the gain of having newly sawn end faces. It is still very likely that newly sawn end faces would give better results than the untreated, if the method was tuned differently.

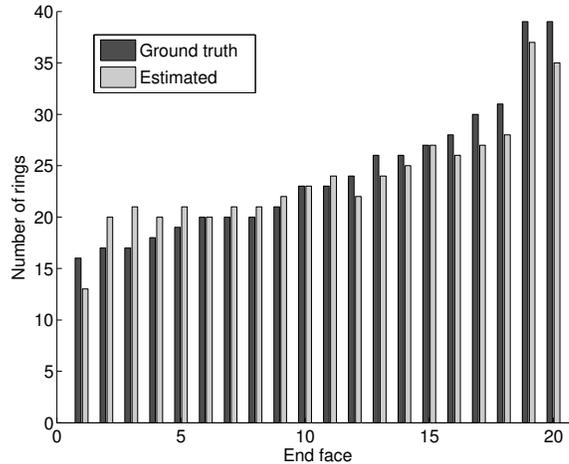


Figure 69: The results for 20 of the images in the Heby2007 dataset.

By examining the results a bit further it was concluded in Paper VII that more than approximately 25 annual rings in the measurement area were impossible to count with the automatic method, since the resolution was too low. It was still possible though, to classify these end faces into the correct quality class.

8.2.1 Additional experiments

The method presented in Paper VI and there used for pine end faces has here been applied to the spruce images in dataset Heby2007. The major difference between the NybyOnline dataset and the Heby2007 set is the image quality regarding noise and the contrast and colour difference between earlywood and latewood. The image quality of the Heby2007 set is better than the NybyOnline dataset. However, many of the end faces have very rough surfaces. The roughness is a difficult disturbance when estimating the number of annual rings, especially in combination with a small colour difference between earlywood and latewood. The annual rings could therefore not be estimated with satisfactory results on all images. As a proof of concept, 20 end faces from the Heby2007 dataset were selected and analyzed. In the selection, end faces with too rough a surface were excluded. For the chosen images, the annual ring pattern was analyzed in one automatically chosen direction and the result was compared to the number of rings counted manually by the author. This is the same approach as was taken also in Paper VI, for the first attempt to count annual rings on pine end faces.

The results from the measurements are shown in Figure 69, where the ground truth denotes the number of rings counted manually along the chosen direction. The results show that it is possible to count the number of rings rather well up to a much higher limit than what was the case in the NybyOnline set, presented in Paper VII. The difference between images that were included and excluded from these analyzed images is illustrated in Figure 70, where one included and one excluded image is shown.

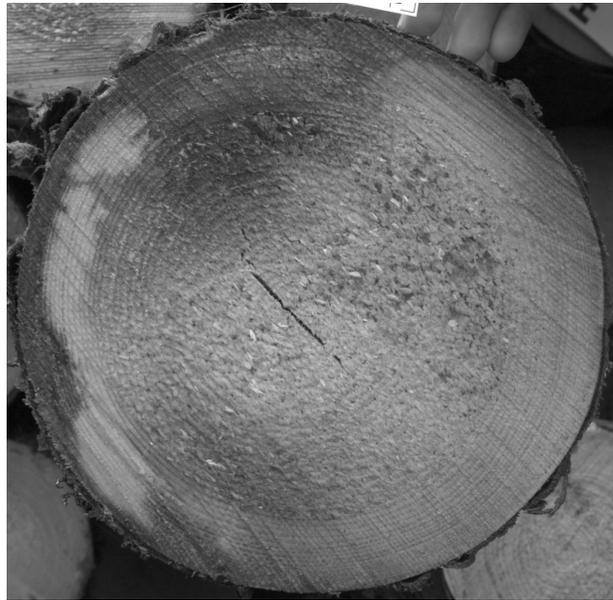


Figure 70: Example of an image that was included in the analyzed dataset (upper) and one that was not (lower).

9 Concluding remarks

In this chapter conclusions are drawn about the work and the developed methods. Ideas for future work in connection with this dissertation are also presented.

9.1 Conclusions

In this dissertation, image analysis methods for automatic quality measurements of log end faces in the sawmill industry were developed and evaluated. The work includes both theoretical tools and practical applications, from image acquisition to measuring end face features. The main contributions are the algorithm for simulating synthetic wood, the theoretical work in connection with the polar distance transform, and the methods for estimating the pith position and the number of annual rings. Furthermore, methods for preprocessing the large amount of data were developed, which were included in the dissertation for completion.

The synthetic end face images were created with the aim to be used in development and evaluation of methods for measuring end face features. A real dataset of approximately 200 images depicting pine was used to determine features such as size, presence of heartwood, number of knots and cracks, and colour, which were then used to create realistic synthetic end faces. Also features related to harvesting, storage of the wood, and image acquisition were simulated. By changing parameters in each of the different steps, different features can be analyzed separately. In the dissertation it was shown that the method to create synthetic data can be extended to 3D, to simulate realistic logs. Unfortunately, time flies and therefore there was no such utilization of the data within the scope of this work.

As a step in the development of the method for annual ring measurements a new distance transform was presented. By computing the radius and angle for a pixel compared to an image origin, distance can be measured in terms of polar coordinates, i.e., the polar distance transform. A grey weighted version was also developed, and in connection with that, a method to compute grey weighted distance in a 5×5 neighbourhood was proposed.

To apply the polar distance transform, the position of the centre of the annual rings, the pith, must be known. An automatic method to estimate the pith position using local orientations in the image was presented in this work. It is robust to noisy and rough end faces as well as disturbances such as knots, cracks, or partial coverage of dirt or snow. This method was evaluated on over 500 images of spruce and pine and used as a starting point for the annual ring measurements.

In the method for measuring the number of annual rings on an end face the circular shape of the polar distance transform is combined with the intensity of the annual ring pattern. A dataset of pine images acquired in sawmill production was used to evaluate the annual ring measurements. The automatic method performed as well as the manual graders when classifying logs into different quality classes based on the number of annual rings. The exact number of rings could not be estimated above a certain number, which was concluded to depend on the resolution of the

images, but the grading was still successful. Initial test have also been performed on spruce end faces which showed promising results for the small set of images. The method for determining the number of annual rings is completely automatic, including cropping the images, filtering to remove marks from uneven sawing, estimating the pith position and measuring the number of annual rings.

For the analysis to be applicable at a sawmill measurement station, it must be carried out during a couple of seconds. No effort was made in this work to optimize methods for this purpose, and therefore no thorough analysis of the required time. However, the choice of methods has been made with computational time in mind. The pith detection method uses filters separable in the spatial dimensions which is effective. The step of combining the orientations into an intersection image is dependent on the size of the image, and the number of orientations taken into account. When it comes to the computations of annual rings the previously computed local orientations are used to determine which areas to measure in, and the polar distance transform is applied only to small parts of the image.

Furthermore, with another camera setup, imaging the end faces from the front at a specific distance, the computations of pixel size is not required, nor is the automatic cropping if the imaged end face is positioned approximately in the centre of the image. When it comes to the method for reducing the marks from uneven sawing, the fast Fourier transform is applied to transform the image to the Fourier domain and back. This can possibly also be avoided since one important reason for the marks is that the illumination creates shadows on the end face. With a more even illumination it is possible that the disturbances are reduced to a level where they are not disturbing to the rest of the measurements.

9.2 Future work

The work in this dissertation includes many different parts, which could all be developed or evaluated further.

- The simulated images resemble pine end faces. However, a few features can be improved to make them even more realistic, e.g., the annual ring pattern and the knots. The annual rings are at present simulated using a sine function. This is not ideal for modelling the growth, therefore a more suitable function would make the result even better. The knots on simulated 2D end faces do not scale with the distance to the pith, in the present model. However, in the 3D model for simulating the knots in logs, cone shaped branches are used, automatically giving this scaling. This should be used also to the 2D case for more realistic knots. A more realistic colour model for the knots is also useful, possibly built up as a growth from the centre of the knot and outwards, simulating the growth of the branch. Simulation of other species could also be performed using the proposed method. A test set of the relevant species is then required to set the parameters, such as colour, presence of heartwood, and presence of knots, to realistic values. Furthermore, a more thorough investigation of what happens in a 3D log needs to be performed so that parameters can be tuned towards real logs. This can be accomplished using the data in the

Swedish stem bank (Grundberg et al., 1995), where data from 200 pine stems from different parts of Sweden are stored. The stems were scanned in a CT-scanner and features like shape, pith position, heartwood size and presence of knots have been extracted.

- When it comes to the polar distance transform, it would be interesting to compare the two methods for computing it, chamfering and fast marching, both regarding computational efficiency and the actual result. To implement the polar distance transform as a weighted distance transform with Dijkstra's wave front propagation method (Dijkstra, 1959) could also be investigated. The results are expected to be equal to the ones in chamfering, but since every pixel is only visited once, the computations are assumed to be more efficient.
- The method for measuring the number of annual rings was thoroughly evaluated for pine end faces in sawmill online production. However, the experiments are limited when it comes to spruce end faces. The small test set gave good results using the automatic analysis but further evaluations of annual ring measurements on spruce need to be performed, possibly with a fixed camera setup and illumination. Then it is possible to obtain a representative dataset regarding roughness of the end faces and contrast between earlywood and latewood, which were the most important differences between spruce and pine observed in this work. If large differences are the case, the method for annual ring measurements might need modifications to handle spruce end faces.
- The synthetic data can be used to analyze the limitations of the methods developed for pith detection and annual ring measurements. By running the methods on different synthetic images, changing one feature at the time for the synthetic data, it is possible to investigate which parts of the methods that need improvement. In this way, the methods can be developed further.

Some work is still needed to take the step from the developed methods presented here to a product working in sawmill environment. Time is an important issue, and the methods must be implemented efficiently, to be used online. A camera setup that is more ideal for the task is also desirable, where the camera can be placed so that it is possible to use flashlight in the imaging, and preferably also so that the end face is imaged directly from the front, without perspective distortion.

Summary in Swedish

Automatisk analys av ändytebilder av stockar i sågverksindustrin

Arbetet i den här avhandlingen har gjorts i samarbete med SDC - skogsnäringens IT-företag, med syftet att ta fram metoder för automatisk mätning av träkvalitet på stockar i sågverksindustrin.

Ändyteanalys görs idag på svenska sågverk som en del av kvalitetsbestämningen av en trästock. Denna analys görs av en virkesmätare genom visuell granskning när stocken passerar på ett transportband genom sågverkets mätstation. I virkesmätarens uppgifter ingår att ta hänsyn till information om stockens yttre dimension, vilken ges av en 3D-mätarm, eventuella skador från fällning, krök på stocken, röta och tjurved samt antalet årsringar på en specifik del av ändytan. Antalet årsringar ska uppskattas till en av tre olika kvalitetsklasser för tall, och en av två olika klasser för gran. Mätningen görs under tiden stocken passerar på transportbandet, vilket är ca två sekunder. Den här avhandlingen utreder möjligheterna för att utföra delar av ändyteanalysen automatiskt med hjälp av bildanalys.

En stor del av den forskning som tidigare har gjorts när det gäller bildanalys av ändytor har gjorts på putsade ytor avbildade i scanners eller genom mikroskop, eller med andra typer av bildtagningar, t ex datortomografi. Mycket av forskningen gäller tillämpningar inom dendrokronologi, läran om åldersbestämning av trä med hjälp av årsringar, där årsringstäthet inklusive tjocklek av vårved och sommarved mäts noga. Även mer sågverksnära tillämpningar har analyserats, men då mestadels genom fotografering i labmiljö av ändytor som har varit antingen behandlade (putsade och oljade) eller nyligen sågade. Hur ändytan har behandlats, i kombination med bildtagningsmetodiken, avgör hur bra läsbarheten i bilden blir. Det gör att olika typer av bildanalysmetoder kan behövas, även om det är samma egenskaper som ska analyseras.

I arbetet till den här avhandlingen används bilder av ändytor som de ser ut i sågverksmiljö. Ett stort antal av bilderna har samlats in under ett års tid med en fastmonterad ändytekamera på Setra Group Nyby sågverk i Björklinge i Uppland. Bilder har tagits varje månad under ett års tid, vilket har resulterat i ändytor från olika årstider, och därmed med olika egenskaper och svårigheter. Dessa data antas därför vara representativa när det gäller förekomst av olika egenskaper hos ändytorna, i alla fall på det aktuella sågverket. Kameran var monterad ovanför transportbandet på mätstationen, vilket gör att bildtagningen är realistisk för den här typen av bilder. Övriga bilder har tagits av ändytor på sågverk eller från stockstumpar samlade efter avverkning. Bland de fotade ändytorna finns många olika egenskaper representerade, såsom kärmved, kvist, tjurved, röta, blåträ, sprickor, ocentrerad mäg och väl synliga möster efter sågningen.

En svårighet med bildanalys är att veta vad det korrekta resultatet är för de mätningar som utförs. Ofta krävs det att en expert på området, i det här fallet trä, noggrant märker ut i bilder eller i riktiga data vad som är korrekt. Även i de fall där

märkningen inte måste utföras av en expert är den ofta tidskrävande och inte alltid entydig, varför detta resultat ibland saknas i bildanalyssammanhang. En lösning till det problemet är att skapa syntetiska data där man på förhand känner till svaret. Om syntetiska data är realistiska kan de användas för att utveckla och utvärdera bildanalysmetoder. I det här arbetet har en generator för syntetiska ändytebilder utvecklats. Genom att analysera riktiga data har de olika egenskaperna anpassats till att efterlikna ändytebilder av tall. I simulatören efterliknas både träet och dess egenskaper, liksom effekter av sågning, lagring, kamerauppställning och bildtagning. I avhandlingen presenteras en utvidgning av metoden till att även simulera tredimensionella stockar.

Avhandlingen innehåller både teoretiska verktyg som är lämpliga för att hantera de riktiga bilderna såväl som mer praktiskt inriktade metoder. Den polära avståndstransformen har presenterats i arbetet och är ett sätt att mäta avstånd där en positions radie och vinkel jämfört med ett origo tas med i beräkningen. Genom att vikta olika riktningar på olika sätt går det att reglera avståndsmåttet så att den kortaste vägen mellan två punkter på samma radie är ett cirkelsegment. Genom att kombinera detta mått med intensiteten hos årsringsmönstret kan antalet årsringar beräknas även i bilder som är brusiga och har ett otydligt mönster.

En förutsättning för att kunna applicera den polära avståndstransformen är att veta positionen för centrum av årsringarna, mörgen. Mörgepositionen kan också vara intressant i sig, eftersom en väldigt ocentrerad mörg kan vara ett mått på tjurved i stocken, vilket bör undvikas i sågade varor. En metod för att hitta mörgen i de aktuella bilderna har utvecklats och utvärderas på ungefär 600 bilder. Metoden baseras på att analysera riktningar i bilden för att på det viset använda årsringsmönstret för att hitta mitten. Genom att mäta var de flesta av riktningarna sammanstrålar går det att estimerar var mörgen finns. I metoden beräknas också ett mått för hur säkert estimatet är. För de ändytor där den automatiska mörgdetektionen är osäker ges därmed en varning till användaren. En osäker mätning kan t ex bero på att ändytan är smutsig eller skadad, eller att bilden är oskarp så att årsringsmönstret är otydligt.

Många ändytor har mönster efter sågningen som uppträder som ränder över ändytan. Dessa ränder stör både mörgestimering och årsringsmätning eftersom de ger upphov till kanter och mönster som kan förväxlas med årsringarnas riktningar. En automatisk metod för att ta bort mönstret som uppkommer vid sågning har också utvecklats i det här arbetet.

Den automatiska årsringsmätningen är utvärderad på ändytebilder av tall som var fotade på en mätstation på ett svenskt sågverk och noggrant kvalitetsbestämt av en kontrollmätare. Försöket visar att resultaten är likvärdiga med resultaten från en manuell mätning vid mätstation, som är standard idag, vilket innebär att metoderna kan användas som en del av en kvalitetsbestämning av sågtimmer, under överinseende av en virkesmätare.

References

- Andreu, J.-P., Rinnhofer, A. (2001). Automatic detection of pith and annual rings on industrial computed tomography log images. In *Proceedings of ScanTech, the ninth international conference on scanning technology and process optimization for the wood industry*. pp. 37–47.
- Andreu, J.-P., Rinnhofer, A. (2003). Modeling knot geometry in Norway spruce from industrial CT images. In *Proceedings of the 13th Scandinavian Conference on Image Analysis (SCIA)*. pp. 786–791.
- Anonymous *PixeLINK* (2004). *PixeLINK-A780 6.6 Megapixel Camera System Guide*. Document No: 04646-01.
- Anonymous *VMR* (2007). *Measuring rules for pine and spruce sawlogs, VMR circular no 1-07*.
- Bærentzen, J. A. (2001). On the implementation of fast marching methods for 3D lattices. Tech. rep., Informatics and Mathematical Modelling, Technical University of Denmark, DTU.
- Bayer, B. E. (1976). Color imaging array. U.S. Patent No. 3,971,065.
- Bhandarkar, S., Faust, T. D., Tang, M. (1996). A system for detection of internal log defects by computer analysis of axial CT images. In *Proceedings of the third IEEE Workshop on Applications of Computer Vision (WAVC)*. pp. 258–263.
- Bigün, J. (1992). Frequency and orientation sensitive texture measures using linear symmetry. *Signal Processing* 29:1–16.
- Blinn, J. F. (1978). Simulation of wrinkled surfaces. In *Proceedings of SIGGRAPH*. 12(3), pp. 286–292.
- Borgefors, G. (1986). Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing* 34:344–371.
- Bowyer, J., Shmulsky, R., Haygreen, J. G. (2007). *Forest products and wood science: an introduction*, 5th Edition. Blackwell, Ames Iowa.
- Buchanan, J. W. (1998). Simulating wood using a voxel approach. *Computer Graphics Forum* 17(3):105–112.
- Canny, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8(6):679–698.
- Cerda, M., Hirschfeld-Kahler, N., Mery, D. (2007). Robust tree-ring detection. In *Proceedings of Pacific-Rim Symposium on Image and Video Technology PSIVT*. pp. 575–585.
- Chalifour, A., Nouboud, F., Deprost, B., Okana, S. (2001). Automatic detection of tree-rings on wood disc images. In *Proceedings of the 5th International Conference on Quality Control by Artificial Vision (QCAV)*. 2, pp. 348–352.

- Conner, W., Schowengerdt, R., Munro, M., Hughes, M. (1998). Design of a computer vision based tree ring dating system. In *Proceedings of IEEE Southwest Symposium on Image Analysis and Interpretation*. pp. 256–261.
- Cooley, J., Tukey, J. (1965). An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation* 19(90):297–301.
- Dijkstra, E. W. (1959). A Note on Two Problems in Connection with Graphs. *Numerical Mathematics* 1:269–271.
- Ebert, D. S., Musgrave, F. K., Peachey, D., Perlin, K., Worley, S. (2003). *Texturing and modeling: a procedural approach*, 5th Edition. Elsevier Science.
- Entacher, K., Hegenbart, S., Kerschbaumer, J., Lenz, C., Planitzer, D., Seidel, M., Uhl, A., Weiglmaier, R. (2008). Pith detection on CT-cross-section images of logs: An experimental comparison. In *Proceedings of the 3rd International Symposium on Communications, Control and Signal Processing (ISCCSP)*. pp. 478–483.
- Entacher, K., Planitzer, D., Uhl, A. (2007). Towards an automated generation of tree ring profiles from CT-images. In *Proceedings of the 5th International Symposium on Image and Signal Processing and Analysis (ISPA)*. pp. 174–179.
- Flood, K. (2003). Segmentation of knots - a problem in 3D X-ray scanning of logs. Ph.D. thesis, Linköping university.
- Fouard, C., Gedda, M. (2006). An objective comparison between gray weighted distance transforms and weighted distance transforms on curved spaces. In *Proceedings of the 13th International Conference on Discrete Geometry for Computer Imagery (DGCI)*. pp. 259–270.
- Gjerdrum, P., Høibø, O. (2004). Heartwood detection in Scots pine by means of heat-sensitive infrared images. *European Journal of Wood and Wood Products* 62(2):131–136.
- Gonzalez, R. C., Woods, R. E. (2002). *Digital Image Processing*, 2nd Edition. Prentice Hall.
- Granlund, G. H., Knutsson, H. (1995). *Signal Processing for Computer Vision*. Kluwer Academic Publishers.
- Grundberg, S., Grönlund, A., Grönlund, U. (1995). The Swedsh stem bank: a database for different silvicultural and wood properties. Tech. rep., Division for wood technology, Luleå university of technology.
- Hanning, T., Kickingereder, R., Casasent, D. (2003). Determining the average annual ring width on the front side of lumber. In *Proceedings of SPIE: Optical Measurement Systems for Industrial Inspection*. Vol. 5144. pp. 707–716.
- Hassouna, M., Farag, A. (2007). Multistencils fast marching methods: A highly accurate solution to the Eikonal equation on cartesian domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(9):1563–1574.

- He, Z., Munro, M. A. R., Gopalan, G., Kulkarni, V., Schowengerdt, R. A., Hughes, M. K. (2008). System and algorithm design for a new generation tree-ring image analysis system. *Optical Engineering* 47(2).
- Heijmans, H. J. A. M., Ronse, C. (1990). The algebraic basis of mathematical morphology. I. dilations and erosions. *Computer Vision, Graphics, and Image Processing* 50(3):245–295.
- Ikonen, L. (2006). Distance transforms on gray-level surfaces. Ph.D. thesis, Lappeenranta University of Technology.
- Jonsson, L. (1992). Registrering av årsringsutveckling med bildanalys. Tech. rep., Swedish university of agricultural sciences, Dept. of forest products. In Swedish.
- Kärenlampi, P. P., Riekkinen, M. (2002). Pine heartwood formation as a maturation phenomenon. *Journal of Wood Science* 48(6):467–472.
- Laggoune, H., Sarifuddin, M., Guesdon, V. (2005). Tree ring analysis. In *Proceedings of Canadian Conference on Electrical and Computer Engineering*. pp. 1574–1577.
- Lefebvre, L., Poulin, P. (2000). Analysis and synthesis of structural textures. In *Proceedings of Graphics Interface*. pp. 77–86.
- Levi, G., Montanari, U. (1970). A grey-weighted skeleton. *Information and Control* 17:62–91.
- Longuetaud, F., Leban, J.-M., Mothe, F., Kerrien, E., Berger, M.-O. (2004). Automatic detection of pith on CT images of spruce logs. *Computers and Electronics in Agriculture* 44:107–119.
- Morales, S., Guesalaga, A., Fernández, M. P., Guarini, M., Irarrázaval, P. (2004). Computer reconstruction of pine growth rings using MRI. *Magnetic Resonance Imaging* 22(3):403 – 412.
- Nyström, J., Hagman, O. (1999). Real-time spectral classification of compression wood in *Picea abies*. *Journal of Wood Science* 45(1):30– 37.
- Peachey, D. R. (1985). Solid texturing of complex surfaces. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques, SIGGRAPH. ACM*, pp. 279–286.
- Perlin, K. (1985). An image synthesizer. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques, SIGGRAPH. ACM*, pp. 287–296.
- Ramanath, R., Snyder, W. E., Bilbro, G. L., III, W. A. S. (2002). Demosaicking methods for Bayer color arrays. *Journal of Electronic Imaging* 11(3):306–315.
- Rosenfeld, A., Pfaltz, J. (1966). Sequential operations in digital picture processing. *Journal of the Association for Computing Machinery* 13(4):471–494.

- Rutovitz, D. (1968). Datastructures for operations on digital images. In *Proceedings of Pictorial Pattern Recognition*. Thompson, Washington, pp. 105–133.
- Serra, J. (1982). *Image Analysis and Mathematical Morphology*. Vol. 1. Academic Press.
- Sethian, J. A. (1999). *Level Set Methods and Fast Marching Methods*. Cambridge University Press.
- Shih, F. Y., Cheng, S. (2004). Adaptive mathematical morphology for edge linking. *Information Sciences—Informatics and Computer Science* 167(1-4):9–21.
- Sjöberg, P.-J., Danielsson, P.-E., Seger, M. M. (2001). Image analysis of annual ring pattern for prediction of wood quality. In *Proceedings of Swedish Symposium on Image Analysis (SSBA)*. pp. 61–64.
- Soille, P. (1999). *Morphological Image Analysis, Principles and Applications*. Springer-Verlag.
- Soille, P., Misson, L. (2001). Tree ring area measurements using morphological image analysis. *Canadian Journal of Forest Research* 31:1074–1083.
- Sonka, M., Hlavac, V., Boyle, R. (1999). *Image Processing, Analysis, and Machine Vision*. Brooks/Cole Publishing Company.
- Toivanen, P. J. (1996). New geodesic distance transforms for gray-scale images. *Pattern Recognition Letters* 17:437–450.
- Tsai, R. Y. (1986). A efficient and accurate camera calibration technique for 3D machine vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 364–374.
- Tsui, K. K. (1996). Using geometric processing in the visualization of ring features in tomographic images of wood. In *Proceedings of the 3rd International Conference on Image Processing (ICIP)*. pp. 547–550.
- Unser, M. (2000). Sampling - 50 years after Shannon. In *Proceedings of the IEEE*. pp. 569–587.
- Wang, X. (1998). Log classification by single X-ray scans using texture features from growth rings. In *Proceedings of the 3rd Asian Conference on Computer vision (ACCV)*.
- Österberg, P. (2009). Wood quality and geometric measurements based on cross section images. Ph.D. thesis, Tampere university of technology.
- Österberg, P., Ihalainen, H., Ritala, R. (2004). Method for analyzing and classifying wood quality through local 2D-spectrum of digital log end images. In *Proceedings of International Conference on Advanced Optical Diagnostics in Fluids, Solids and Combustion*.

Acknowledgements

I would like to express my sincere gratitude to a number of people who supported me during my work.

- © First of all, my main supervisor Gunilla Borgefors, who guided me during my doctoral studies with many good advice and encouraging words. My assistant supervisor, Lars Björklund for providing knowledge and insight in the procedure of measuring sawlogs, and for making it possible to cooperate with sawmill industry. My assistant supervisor Mats Nylinder, our meetings have always brought a lot of inspiration.
- © All people involved in the collection of data and image acquisition: Fredrik Nylund, Magnus Jansson, Maria Jonsson, Hans Fryk, Ari Karjula, Christina Lundgren, Per Berg, Jonas Lantz, Jan Olsson, Niklas Anderén, Christer Forsmark, and Jacob Edlund.
- © Ewert Bengtsson, and all seniors at CBA for creating an inspiring research environment. A special thanks to Stina Svensson and Joakim Lindblad for supervision and collaboration, mainly during my second year of PhD studies. Robin Strand, for being my co-author and for always having time for discussions. Catherine Östlund for involvement in my work, especially when it comes to the practical issues with the camera setup. Ingela Nyström for listening, for the support you always give, and for good friendship.
- © Lena Wadelius, for doing all the things at CBA that you seldom get credit for. Your work is much appreciated. Olle Eriksson, for taking care of my computer and everything connected to it, and for your impressively short response time. Bo Nordin and Eva Pärt-Enander, for inspiring my teaching in such different ways.
- © All fellow PhD-students and colleagues at CBA during this time, especially my roommates: Ida-Maria, Hamid, Erik, Mathias, Khalid, Filip and Magnus. The days at CBA would not have been as fruitful or fun without you!
- © Nina Kingstad, for all the fun we have had knitting, shopping and drinking tea.
- © My family for believing that I am capable of doing anything. Your support has meant a lot to me during all my studies.
- © Björn, thank you for always being there, for being calm when I am not, and for making me braver than I otherwise would be.

The financial support from SDC - skogsnäringens IT-företag, is gratefully acknowledged.

Errata to Paper III

In Paper III the weights c_1 and c_2 should be replaced by $\frac{1}{c_1}$ and $\frac{1}{c_2}$, respectively in Equation 3. In Equation 8 and the following expressions for the weights c'_1 and c'_2 , also these should be replaced by $\frac{1}{c'_1}$ and $\frac{1}{c'_2}$, respectively. The correct equations are

$$\left(\frac{1}{c_1}D_{\mathbf{v}_1}T(a,b)\right)^2 + \left(\frac{1}{c_2}D_{\mathbf{v}_2}T(a,b)\right)^2 = \frac{1}{F(a,b)^2}, \quad (3)$$

$$\left(\frac{1}{c'_1}\frac{\partial T}{\partial x}\right)^2 + \left(\frac{1}{c'_2}\frac{\partial T}{\partial y}\right)^2 = \frac{1}{F(a,b)^2}, \quad (8)$$

$$\frac{1}{c'_1} = \frac{1}{c_1}\left(1 - \frac{\theta}{\pi/2}\right) + \frac{1}{c_2}\left(\frac{\theta}{\pi/2}\right)$$

$$\frac{1}{c'_2} = \frac{1}{c_1}\left(\frac{\theta}{\pi/2}\right) + \frac{1}{c_2}\left(1 - \frac{\theta}{\pi/2}\right).$$

The computations in the paper are correct. For the text and the figures, the weights are correct. A high weight gives slow propagation in the corresponding direction.