Improving knot segmentation using Deep Learning techniques

Stefano Giovannini
Department of Information Engineering (DEI), Università di Padova
Davide Boschetto
Microtec GmbH, Venice, Italy, davide.boschetto@microtec.eu
Enrico Vicario
Microtec GmbH, Venice, Italy, enrico.vicario@microtec.eu
Mauro Cossi
Microtec GmbH, Venice, Italy, mauro.cossi@microtec.eu
Andrea Busatto
Microtec GmbH, Venice, Italy, andrea.busatto@microtec.eu
Stefano Ghidoni
Department of Information Engineering (DEI), Università di Padova, ghidoni@dei.unipd.it
Enrico Ursella *
Microtec GmbH, Venice, Italy, enrico.ursella@microtec.eu
Dipartimento di Scienze Ambientali, Informatica e Statistica, Università Ca' Foscari Venezia

* Corresponding author

Abstract
In the context of Computed Tomography scanning of logs, accurate detection of knots is key for delivering a successful product. Reliable detection of knots in the sapwood is hard with traditional computer vision techniques, because of the different density conditions between sapwood and heartwood. The advancements provided by deep learning in the field of semantic image segmentation kick-started a new way of approaching such problems: deep neural networks can be trained on large amounts of labelled data and successfully employed in production environments to improve the performances on knot detection. Adapting state-of-the-art network architectures and using more than 10.000 labelled knots from pine and spruce logs, we were able to develop a new two-step approach for identifying knots in CT scans of logs with unprecedented accuracy while at the same time satisfying the time constraints that a real-time industrial application needs. The first step runs on the log's axis, while the second runs on each candidate knot's axis. False positives from the first step are very rare (even with dry/dried logs), so no computational power is wasted for non-existing knots. Using this approach, we are able to see the internal defects of a log in real time in the production chain without having to cut it first, therefore being able to optimize even more the output of the chain on each client's requirements.

Keywords: CNN, U-Net, semantic segmentation, knot detection, CT, dead knot border

Introduction
Computed Tomography (CT) for production optimization in sawmills has been a reality for a few years. CT Log (Giudiceandrea et al., 2011) is a CT scanner produced by Microtec able to measure logs at up to 180m/min, compute a model of the internal features of each log and optimize the entire process based on the characteristics of the raw material and the production requests of the sawmill. The steps of bucking, sorting and optimizing the cutting pattern can be improved with the use of the internal information provided by computed tomography. Many works demonstrated the economic advantage of optimizing the sawing process with the use of CT images (Rais et al 2017) (Berglund et al 2013) (Stangle et al 2015).

The automatic detection of the internal features of a log from CT images has been addressed by many works, especially for knots detection (Andreu at al 2003) (Breinig et al., 2012) (Fredriksson et al., 2017) (Cool et al., 2017) (Longuetaud et al., 2012). Only a few of them (Oja, 2000) (Johansson et al.,

2013) addressed the problem of the detection of the dead knot border. The dead knot border is the point that divides the part of knot that is sound from the one where the knot is dead. On the dead part of a knot, a thin bark layer divides the knot from the rest of the wood: therefore, the mechanical connection with the rest of the wood is lower, sometimes causing the knot to fall off the board. For this reason a clear estimation of the point where a knot becomes dead is very important in order to be able to optimize the cutting pattern and produce boards with higher quality. In (Oja, 2000) the coefficient of determination of the linear regression between the predicted and measured percentage of sound knots on each board was measured to be R^2=0.72. In (Johansson et al., 2013) the detection was accomplished measuring the point where the diameter of each knot stops growing. The RMSE of the dead knot border estimation on Pine logs was of 12mm.

The aim of this work is presenting how Convolutional Neural Networks (CNN) were applied in order to improve the detection of knots from CT images. The detection was performed in two steps: the first step applies a semantic segmentation on the whole log in order to define the position of each knot. In the second step, an area around each knot is analyzed in order to calculate its properties, and in particular the measurement of the dead knot border.

The neural network in charge of the semantic segmentation is a fully convolutional network that performs 2D convolutions on volumes of consecutive slices of CT images to produce probability maps expressing the likelihood that each pixel is part of a knot. The second network has the purpose of classifying volumes of knots as sound or dead and of identifying the right dead knot border position.

One of the important requirements of using deep learning is that a big number of samples must be collected and labelled with accurate information to correctly train the system. The correct labelling requires a lot of work from trained people but also a correct procedure. For this reason, a specific software was developed for labelling the dataset by visual inspection of CT images. The definition of the dead knot border from CT images was a hard task for our graders, therefore we chose to use measurements taken directly on the surface of the boards after the logs were sawn, since the status of a knot is clearly visible on the surface of a board but not as clearly interpretable from CT images.

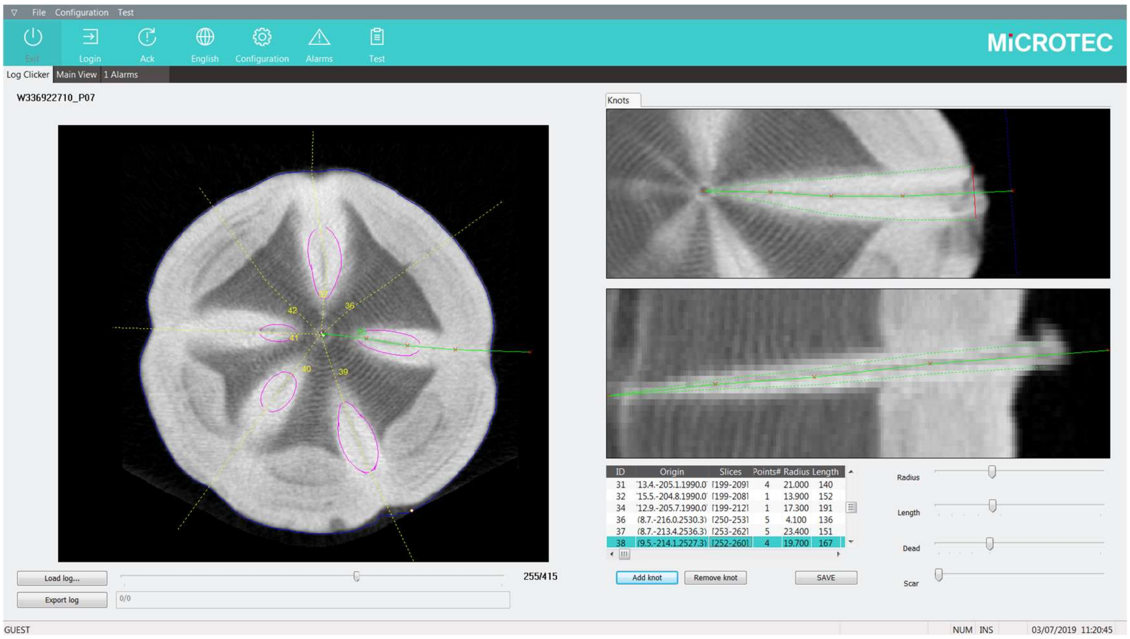

Figure 1— Visualization of the software for the manual labelling of the knots. Different views of the same area are shown in order to help a precise definition of the positions.

Material and methods

Step 1: Knot identification

CT Log produces 3D images where each voxel has the dimension of 1x1 mm in transversal direction and 10 mm along the axis of the log. In the remainder of this paper we will call x and y the first two coordinates of the CT images, transversal to the axis of the log, and z the third coordinate.

In order to visualize and label in 3D each knot, we developed a software shown in Figure 1 where different views of the same part of the log were presented in the same screen in order to label the starting point, end point, dead knot border and diameter profile of the knots. It was also possible to define any number of intermediate points along the trajectory of the axis of the knot.

The CT images of 75 Scots Pine logs (Pinus Sylvestris) that were scanned in different sawmills in Europe were collected in order to create a database. The knots of those logs were manually marked with the software described before and a total of 10.118 knots were collected.

The parametric labelling was transformed to produce a 3D volume of voxels indicating whether each voxel belongs to a knot (voxel value = 1) or is not part of a knot (voxel value = 0). Each slice was scaled in order to have consistent slices dimension (128 x 128 pixels) on all logs. Then, as network input, groups of images composed of 5 adjacent slices were created.

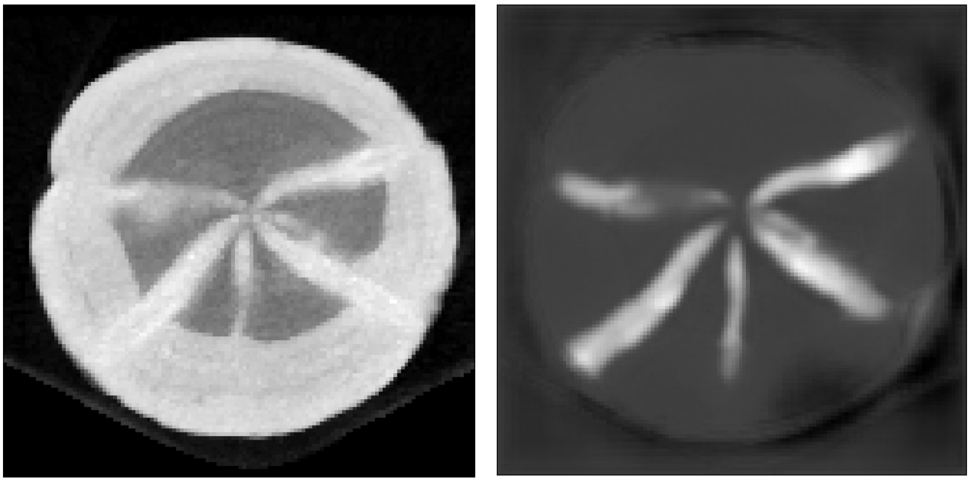

Figure 2 — The original image scaled at 128x128 pixels (left), the CNN segmentation (right)

The CNN was then able to produce a 3D image of the same size of the input were each voxel expresses the probability of it being part of a knot, as shown in Figure 2.

In order to identify the position and the bounding box of each knot, a special version of the Hough transform (Ballard 1981) was implemented. One strong simplification of the model comes from the fact that almost all knots in a log start from the pith, since epicormic knots are very rare in forests. The position of the pith along the log can be easily calculated (e.g., (Boukadida et al., 2012)): a set of values xPith(z) and yPith(z) is therefore obtained. The axis of a knot can then be parametrized with 3 parameters:

- zStart: the z coordinate of the position where the knot starts
- Orientation: the angle of the direction of the knot in the x,y plane;
- Slope: the inclination of the direction of the knot in the z direction with respect to x-y plane.

The algorithm creates a 3D Hough map based on the 3 parameters by looping on a range of possible values of slope between -30% and 30% at step of 2%. Given the slope, for every voxel a unique knot axis passing through both it and the pith exists. So, it is possible to compute the pith position in which the knot starts zStart(x,y,z,slope), where x,y,z are the position of a generic voxel. Then we can calculate the orientation of the knot as orientation(x,y,z)=atan2(yPith(zStart),xPith(zStart)). With this functions it is possible to calculate for every slope and x,y,z voxel, the correspondent zStart,orientation coordinate in the Hough map and add the probability value calculated with the CNN in order to compute the probability of a knot with the given parameters. Choosing the best local maxima of the map allows to create a list of axes of the knots. An example of two layers of a Hough map is shown in Figure 3.

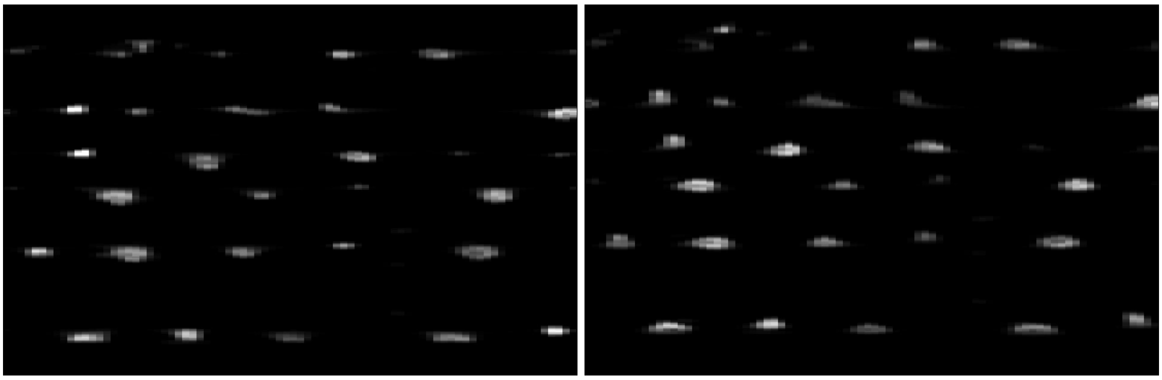

Figure 3—The Hough map of the knots: the orientation in the horizontal axis, zStart in the vertical axis. On the left image the plane with slope=0% is shown, on the right the plane with slope=4%.

Step 2 knot area analysis

Once the axis of each knot has been identified, a 3D volume of voxels is extracted. We define the 3 directions of the extracted boxes as:

- radial direction (r): along the orientation of the knot in the x-y plane
- tangential direction (t) orthogonal to r and z directions
- z direction

The extracted voxels groups are volumes with fixed dimension of 160 x 80 x 80 in the r, t, z directions, respectively. A different scale factor is applied to the three directions in order to optimally fit each knot in a volume, depending on the radius of the log, the maximum expected diameter and the slope. We obtain volumes as in Figure 4 (top).

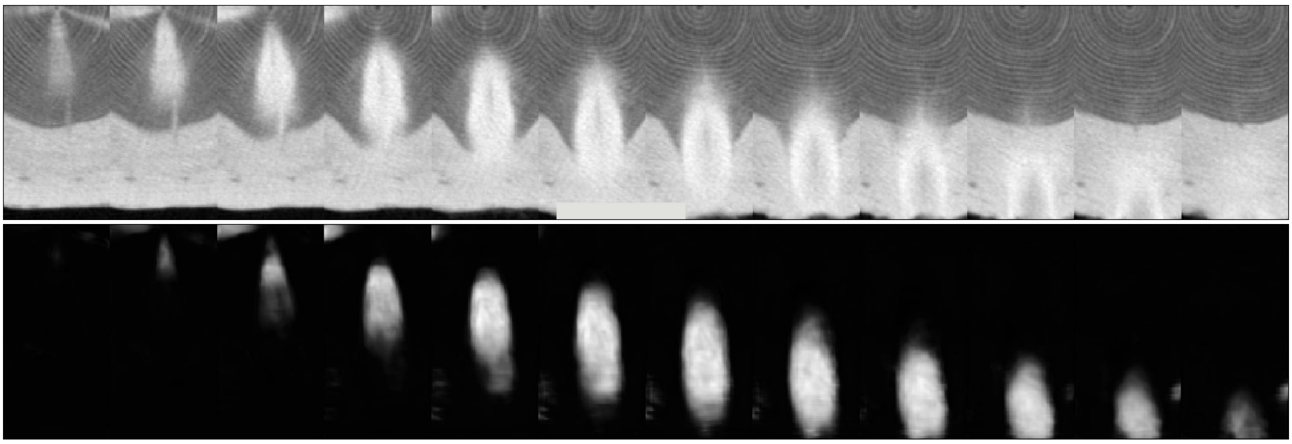

Figure 4 — The block of pixels around a knot. Original image (top) and result of CNN for semantic segmentation (bottom)

One possible solution is to perform a new semantic segmentation on these knot blocks. The analysis of the dimension and center of the segmented voxels along the r direction allowed to compute the direction, dimension and length. The dead knot border is calculated as the point where the dimension of the knot stops growing along the r direction.

In order to improve performances specifically on the detection of the dead knot border and knot diameter, a more reliable ground truth, specific to these two metrics, is needed. The only way to extract reliable ground truth on those metrics was to check the appearance of the knots on board surfaces instead of using CT images. 13 logs of Scots Pine were scanned with a CT Log scanner, sawn in thin boards 15mm thick and the knots were manually measured on the surface of the boards. A reference was taken on the logs by drilling some holes so that it was possible to compute the position of a knot in the CT image given the measured position on the board and vice versa. On the 13 logs, 2412 knots were measured on the surface of the boards.

Figure 5 — An example of reference taken on a log that is CT-scanned and then sawn in boards (left), the manual measurements superimposed on the CT image (right).

In Figure 5 (right) the original CT images are visible with the position of the boards and the measure of the knots taken manually superimposed. In particular, the minimal, maximal diameter, position and dead/sound status were annotated. The information related to each knot was reported in the reference system of the specific block of voxels extracted in order to train a neural network.

The problem is that the requested information (dead knot border and diameter profile) would require a dense ground truth along the radial direction, while only a few manual measurements are available depending on the thickness of the boards and their angle with respect to the knot axis. It was not possible to train a network on the whole knot volume when the ground truth was valid only for a few points along its length (only 1 or 2 points).

To correctly train the network we decided to extract sub-blocks of 11 slices along the r direction around positions where the ground truth was available. Two different networks were trained: one to compute the dead/sound status, the other for the computation of the diameter of the knot.

During the training of the sound/dead network it was possible to extend the ground-truth information also to other points of the knot. If a knot was marked as sound at a certain coordinate r_alive, for obvious biological reasons the knot was alive also at all r<r_alive. For the same reason, if a knot was marked as dead at a certain point, all the subsequent slices have to be marked as dead. This allowed to create a dataset with a high number of samples. At this point the network is able to classify the single slice of a knot to either dead or sound, but it is obviously faster to infer this information by comparing the actual distance of a slice from the pith with the dead knot border value of that knot. To find a

knot’s dead knot boder, for computational time constraints, one slice every six was tested. Once the point where the status passes from sound to dead is found, we then refine the detection to pinpoint the exact slice. After the calculation of the dead knot border we verified the performance of the system by comparing the predicted status with the status of the knots in the sawn boards.

As ground truth for the computation of the diameter we decided to use as training set only the slices of knots where there was a manual measurement. An interpolation of the measured diameter in the slices between two consecutive measurements were possible, but it could have reduced the precision of the system. As long as we can consider that branches are not elliptical, we use the minimal diameter measured on the board as the diameter of the knot in the 3D image.

Results

Knots segmentation

To design and train the networks, we used a computer running Windows 10 Pro, Keras 2.2.4 with Tensorflow 1.13.1 as backend. The first network, aimed at semantic segmentation, has a total of 1.962.913 parameters. It follows the U-Net architecture (Ronneberger et al., 2015) with skip connections and convolutional blocks consisting of 2 consecutive convolutional layers. The first layer interprets the channel axis as a depth axis. Starting from an image size of 128x128 with 5 channels, it compresses the image to a size of 8x8 with 256 channels in the center. Then, in the so-called “upward path” of U-Net, the image is upscaled to the original resolution with 1 output class as channel (the probability of a pixel part of the central slice to be part of a knot). Each convolutional layer applies 3x3 kernels, and for the optimization the Adam optimizer has been used with a learning rate of 0.0002 with binary crossentropy as the default loss function. Early stopping and learning rate reduction on plateau have been used during the training process. The inference time for the computation of a log 4.2 m long was 650 ms. All computation times were measured on a computer using an RTX 2080 GPU on an Intel Core i7-4770 3.4GHz.

Dead knot border

In total the 13 logs presented 634 knots. Each knot intersected one or more boards, the manual measurement were taken at those intersections. The 634 knots intersected the boards in 2412 measured points. 1835 knot intersections were randomly chosen for training and validation set (75% for training and 25% for validation). They generated 34917 knot slices with known dead/sound status and used for training and validation of the neural network. 577 knots intersections, belonging to 158 knots, were used for the test of the detection of the final algorithm estimating the dead knot border.

The test of the performance was done comparing the status of the knots manually measured on the boards with the expected status based on the estimated dead knot border. The results are presented in Table 1.

Table 1—Confusion matrix of the classification of the sound/dead knot classification

		predicted	
		Sound	dead
actual	sound	301 (88.5%)	39 (16.4%)
	dead	39 (11.4%)	198 (83.5%)

The inference of a single slice of the network used for the dead/sound classification required 0.42 ms. Every knot required the inference of 23 knot slices, so in total the computation time for the dead knot border of a knot was 10 ms.

Knot diameter

For the training and validation of the network, 1776 intersections of the knots with boards were used (75% for training and 25% for validation), while 564 were used for the test.

The standard deviation of the difference between the manual measurement and the predicted value of the diameter was 3.2 mm, the average -0.1mm. In Figure 6 a comparison of measured and predicted diameters is shown.

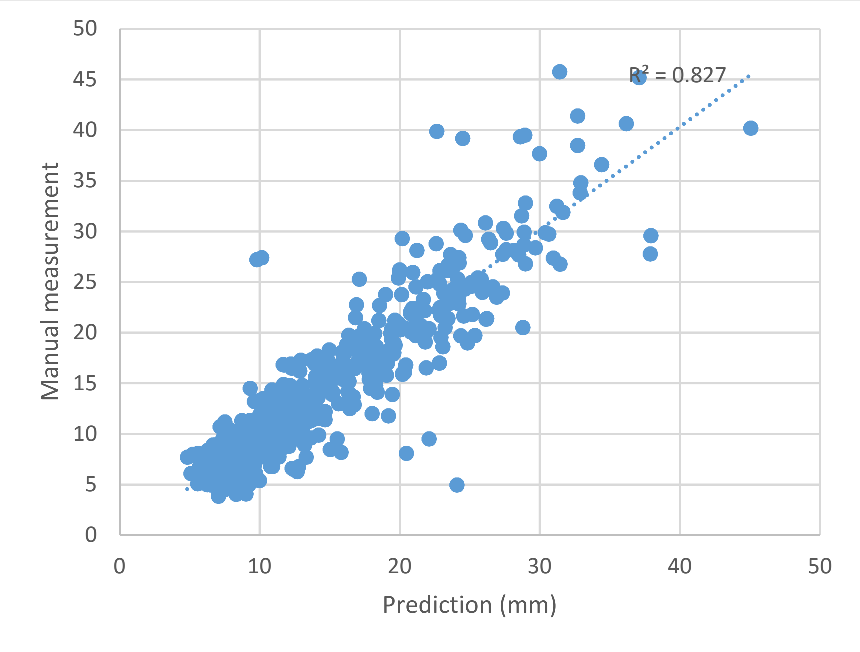

Figure 6— Comparison of the diameter estimation of the knots

The inference of a single slice of the network for the diameter estimation required 0.58 ms. In total, 12 volumes needed to be inferred in order to compute the diameter along each knot, requiring 6.7 ms per knot.

In total, the computation time for a 4.2 m long log with 50 knots was 1450 ms on a single computer with one GPU.

Conclusion and future work

We presented an algorithm that used convolutional neural networks for the automatic detection of the parameters of knots from CT images of Scots Pine logs. In particular we proposed a procedure that allows to train the network with thousands of samples. In particular allowed to train the system using reliable information measured on the surface of boards instead of using CT images. Even if the results of the presented work are very positive, we plan to extend the approach using board scanners installed in a production line to create the ground truth for the training of new networks. Training neural networks on CT images from all logs from a sawmill using as ground truth the scanned data of those logs' sawn boards will create more accurate automatic inspection processes tailored on the specific characteristic of the logs and productions.

References

Andreu, Jean-Philippe, and Alfred Rinnhofer. "Modeling of internal defects in logs for value optimization based on industrial CT scanning." *Fifth International Conference on Image Processing and Scanning of Wood*. Bad Waltersdorf Austria, 2003.

Ballard, D. H. (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern recognition*, *13*(2), 111-122.

Berglund, A., Broman, O., Grönlund, A., & Fredriksson, M. (2013). Improved log rotation using information from a computed tomography scanner. *Computers and electronics in agriculture*, *90*, 152-158.

Boukadida, H., Longuetaud, F., Colin, F., Freyburger, C., Constant, T., Leban, J. M., & Mothe, F. (2012). PithExtract: A robust algorithm for pith detection in computer tomography images of wood–Application to 125 logs from 17 tree species. *Computers and electronics in agriculture*, *85*, 90-98.

Breinig, L., Brüchert, F., Baumgartner, R., & Sauter, U. H. (2012). Measurement of knot width in CT images of Norway spruce (Picea abies [L.] Karst.)—evaluating the accuracy of an image analysis method. *Computers and electronics in agriculture*, *85*, 149-156.

Cool, J., Fredriksson, M., & Avramidis, S. (2017). Knot detection in coarse resolution CT images of logs. In *23rd International Wood Machining Seminar, Warsaw, Poland, 28-31 May 2017*.

Fredriksson, M., Cool, J., Duchesne, I., & Belley, D. (2017). Knot detection in computed tomography images of partially dried jack pine (Pinus banksiana) and white spruce (Picea glauca) logs from a Nelder type plantation. *Canadian Journal of Forest Research*, *47*(7), 910-915.

Giudiceandrea, F., Ursella, E., & Vicario, E. (2011, September). A high speed CT scanner for the sawmill industry. In *Proceedings of the 17th international non destructive testing and evaluation of wood symposium* (pp. 14-16). Sopron, Hungary: University of West Hungary.

Johansson, E., Johansson, D., Skog, J., & Fredriksson, M. (2013). Automated knot detection for high speed computed tomography on Pinus sylvestris L. and Picea abies (L.) Karst. using ellipse fitting in concentric surfaces. *Computers and electronics in agriculture*, *96*, 238-245.

Longuetaud, F., Mothe, F., Kerautret, B., Krähenbühl, A., Hory, L., Leban, J. M., & Debled-Rennesson, I. (2012). Automatic knot detection and measurements from X-ray CT images of wood: a review and validation of an improved algorithm on softwood samples. *Computers and Electronics in Agriculture*, *85*, 77-89.

Oja, J. (2000). Evaluation of knot parameters measured automatically in CT-images of Norway spruce (Picea abies (L.) Karst.). *European Journal of Wood and Wood Products*, *58*(5), 375-379.

Rais, A., Ursella, E., Vicario, E., & Giudiceandrea, F. (2017). The use of the first industrial X-ray CT scanner increases the lumber recovery value: case study on visually strength-graded Douglas-fir timber. *Annals of forest science*, *74*(2), 28.

Ronneberger, O., Fischer, P., & Brox, T. (2015, October). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234-241). Springer, Cham.

Stängle, S. M., Brüchert, F., Heikkila, A., Usenius, T., Usenius, A., & Sauter, U. H. (2015). Potentially increased sawmill yield from hardwoods using X-ray computed tomography for knot detection. *Annals of forest science*, *72*(1), 57-65.