

Article

Performance of Neural Networks in Automated Detection of Wood Features in CT Images

Tomáš Gergel¹ , Ondrej Vacek^{1,2} , Miloš Gejdoš^{1,3,*} , Diana Zraková¹, Peter Balogh¹ and Emil Ješko¹

¹ Forest Research Institute Zvolen, National Forest Centre, T. G. Masaryka 2175/22, SK 960 92 Zvolen, Slovakia; tomas.gergel@nlcsk.org (T.G.); ondrej.vacek@tuzvo.sk (O.V.); diana.zrakova@nlcsk.org (D.Z.); peter.balogh@nlcsk.org (P.B.); emil.jesko@nlcsk.org (E.J.)

² Department of Mathematics and Descriptive Geometry, Faculty of Wood Sciences and Technology, Technical University in Zvolen, T.G. Masaryka 24, SK 960 01 Zvolen, Slovakia

³ Department of Forest Harvesting, Logistics and Ameliorations, Faculty of Forestry, Technical University in Zvolen, T.G. Masaryka 24, SK 960 01 Zvolen, Slovakia

* Correspondence: gejdos@tuzvo.sk; Tel.: +421-455-206-283

Abstract

Computed tomography (CT) enables non-destructive insight into internal log structure, yet fully automated interpretation of CT images remains limited by inconsistent annotations, boundary ambiguity, and insufficient spatial context in 2D slice-based analysis. These challenges restrict the industrial deployment of deep learning for wood quality assessment. This study applies artificial intelligence (AI) and deep learning to the automated analysis of computed tomography (CT) scans of wood logs for detecting internal qualitative features and segmenting bark. Using convolutional neural networks (CNNs), trained models accurately distinguish healthy and damaged regions and segment bark, including discontinuous parts. We introduce a novel pseudo-spatial representation by merging consecutive slices into red–green–blue (RGB) format, which improves prediction accuracy and model robustness across logs. To enhance interpretability, Gradient-weighted Class Activation Mapping (Grad-CAM) highlights regions contributing most to defect detection, particularly knots. Comprehensive evaluation using Sørensen–Dice similarity coefficients and confusion matrices confirms the effectiveness of the proposed approach under industrial conditions. These findings demonstrate that AI-driven CT image analysis can address key limitations of current log-grading workflows and enable more reliable, objective, and scalable quality assessment for timber-dependent economies.

Keywords: neural network; bark detection; CT scan; roundwood; image annotation



Academic Editors: Maria
Francesca Cataldo and Salvatore
F. Papandrea

Received: 2 March 2026

Revised: 18 March 2026

Accepted: 25 March 2026

Published: 27 March 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Wood processing has become increasingly sophisticated, with accurate analysis of raw material quality playing a crucial role [1]. One of the fundamental challenges for modern sawmills is the reliable characterization of internal log structure, since features such as cracks, knots, decay, and deformations directly determine product value and processing strategy [2,3]. However, conventional approaches based on visual inspection or manual interpretation of computed tomography (CT) images remain time-consuming, subjective, and prone to inconsistencies among operators [4]. These limitations highlight the need for objective, scalable, and reproducible evaluation methods.

Significant progress in artificial intelligence, and particularly in deep learning, has opened new opportunities for automating the analysis of image data across industrial

domains [5–7]. Convolutional neural networks, especially convolutional neural networks (CNNs), outperform traditional feature-engineering approaches in a wide spectrum of detection and classification tasks and have demonstrated strong generalization to complex visual patterns. Applying CNN-based analysis to CT scans of logs therefore represents a natural progression toward automated, high-resolution interpretation of internal wood structure, enabling real-time decision-making and reducing human-induced variability [8–11].

Recent research has shown that advanced neural architectures combined with specialized preprocessing pipelines can substantially improve defect detection accuracy in wood. Zhang et al. [6] proposed a hybrid DELM method integrating CNNs with Extreme Learning Machines and NSST-based preprocessing, achieving near real-time performance with a 96.72% detection accuracy. An et al. [12] introduced an enhanced YOLOv8 variant (CWB-YOLOv8) trained on a large multi-species dataset, demonstrating increased robustness to defect variability. Kılıç et al. [13] presented the WD Detector system, using Xception-based feature extraction combined with classical machine-learning classifiers, reaching 99.32% accuracy in surface defect detection. These studies underline a rapid rise in the capabilities of deep learning for wood-quality assessment and motivate its extension to volumetric CT data.

Despite this progress, developing dependable deep learning models for CT log analysis remains challenging—primarily due to issues related to annotation quality, missing standards, data imbalance, and the intrinsic 3D nature of CT imagery. Annotation of internal features is inherently subjective: experts may differ in how they delineate defect boundaries or classify borderline cases, resulting in inconsistent training labels [6]. Human annotation errors—including missed defects, misclassification of subtle features, or partially drawn masks—further complicate model training and can reduce segmentation accuracy [14]. A major obstacle is the absence of unified annotation protocols for wood CT data, unlike the well-established standards in medical imaging [15]. Datasets are often imbalanced, with some defect types underrepresented, which biases learning [6]. Moreover, the three-dimensional complexity of CT volumes requires annotators to consider multiple adjacent slices, making the process slow, labor-intensive, and difficult to quality-control [16].

Addressing these issues calls for improved annotation workflows, standardized labeling practices, multi-expert validation, and the use of automated or semi-automated tools. High-quality annotated datasets are essential for developing reliable deep learning models and represent one of the bottlenecks for industrial deployment.

The present study investigates the use of neural networks for the detection of internal features in CT scans of logs, with the objective of identifying practical solutions applicable under real industrial conditions. Our primary contribution is the introduction and evaluation of a pseudo-spatial RGB representation, in which three consecutive grayscale CT slices are merged into a single color image. This simple but effective approach provides the classifier with additional spatial context without requiring a fully 3D architecture and demonstrably improves robustness and accuracy. A secondary contribution is the enhanced segmentation of bark achieved by enriching the training dataset with challenging and irregular bark cases. Finally, we use Gradient-weighted Class Activation Mapping (Grad-CAM) to interpret model predictions and localize defect-relevant regions, increasing transparency and supporting practical deployment of the method.

Given that many countries—particularly those whose economies strongly depend on timber production—rely on precise log grading to maximize value recovery, the development of reliable automated CT-based defect detection methods has direct practical significance. Improved accuracy and robustness in identifying internal wood features can therefore contribute to more efficient resource utilization, higher product yield, and increased economic competitiveness in timber-oriented markets.

2. Materials and Methods

The training of neural networks was conducted using a continuously developed database of CT scans of log sections (LIBROARBOR) within the LignoSilva Centre of Excellence (National Forest Centre, Zvolen, Slovakia).

The X-ray source is positioned opposite the detectors on both sides of the scanning ring. Computer processing assigns an appropriate radiation absorption value to each point within the cross-sectional plane. Based on these absorption values, a grayscale image is generated. The acquired data were then used to reconstruct slices in the required planes. From CT images—produced as cross-sectional slices composed of voxels measuring 1 mm × 1 mm × 10 mm—a 3D model of each log section was created.

For labeling individual features in CT images, the following methodology was applied. The first feature to be correctly detected is the wood species and the position of the pith in the CT image. This detection was successfully implemented in a previous study [17]. Subsequently, the neural network was trained for automatic bark detection.

Training was performed using annotated CT images in the MATLAB (version R2024a) environment, employing the built-in Image Labeler application. Using manual marking tools, three categories were labeled in each image: background, log (i.e., internal wood mass), and bark.

Annotations were created by a single trained annotator, and the annotated dataset was subsequently reviewed by three independent expert controllers, who visually inspected all labels and corrected any overlooked errors to ensure high annotation quality. Because all annotations were performed by a single annotator, inter-rater agreement could not be evaluated. Metrics such as Cohen's kappa require at least two independent annotators and are therefore not applicable in this study.

These annotations served as input data for training classification models, enabling automatic identification of these regions in new CT images. This procedure ensures high segmentation accuracy and forms the basis for further detection of internal wood defects. Bark annotation was carried out on three CT scans of oak log sections (Figure 1), representing 150 CT images.

To increase robustness and address annotation imbalance, the training database was subsequently supplemented with additional challenging examples—100 CT images containing damaged or discontinuous bark and 100 images with intact bark. Furthermore, 30 CT slices containing knots were included to enrich representation of internal features that may appear near bark boundaries. A total of 380 CT images were therefore used to train the neural network for automatic bark detection in CT scans of logs. A comprehensive summary of the dataset used in this study is presented in Table 1.

After successful training of the neural network for bark detection, the process continued with detecting other internal features such as cracks, decay, and healthy or unhealthy knots. Given the variability of these features, it was necessary first to train the network to distinguish whether a CT image contains a defect or not. This information was recorded in an MS Excel table, where each CT image was assigned a code: absence of defect—0, uncertainty (image excluded from training)—1, and presence of a qualitative feature—2. A custom software interface developed in MATLAB (version R2024a) was used to streamline the labeling and classification of hundreds of CT images. The application is designed for the specific classification of features and displays the probability of defect presence in each CT image. It allows image loading, selection by ID, defect marking, Excel export, and visual verification using a color scale. Functions such as “Record Defects,” “Check,” “Excel with Labels,” and “Select New Excel” simplify the entire annotation and data preparation process for subsequent analysis phases (Figure 2).




Figure 1. Annotation of CT images for training automatic bark detection in MATLAB environment.

Table 1. Dataset summary.

Image Type	Count	Description
Number of scanned logs	3	Selected oak trees, scanned by computed tomography
The number of selected annotated CT scans from three oak logs	150	150 manually annotated slices from three oak logs; additional 100 slices with discontinuous or damaged bark and 100 slices with intact bark
Addition to the training file on the specificity of the CT image	100	Selected CT images from various log scans containing a specific qualitative feature—damaged bark
Addition to the training file on the specificity of the CT image	100	Selected CT images from various log scans containing a specific qualitative feature—intact bark
CT slices containing knots	30	Images with healthy or unhealthy knots used to enrich the segmentation and classification data
Validation set	50	Independently annotated CT slices used for evaluating segmentation accuracy

The developed application also verifies whether the CT image is in grayscale format (and converts it if necessary) and checks whether the image name corresponds to the Excel file name. It enables control of automatic classification for defect presence or absence, including numerical probability values.

Post-processing rule: To correct boundary-related misclassifications during bark segmentation, a simple post-processing rule was applied: pixels originally classified as “Outskirt” with CT intensity above 80 (in the calibrated scale) were reclassified as bark. This rule targets cases where the bark is present but partially detached or low-contrast, resulting in Outskirt—Bark mislabels.



Figure 2. Application interface for simplified annotation of CT images for defect presence or absence.

Neural Network Architecture

Two neural network architectures were used due to the different nature of the tasks. DeepLab v3+ was selected for semantic segmentation of bark because of its strong performance on boundary-sensitive segmentation tasks. In contrast, Inception-v3 was used for defect classification, where global image features are more relevant than precise spatial boundaries. This separation of architectures reflects the distinct requirements of segmentation and classification and avoids unnecessary complexity in model design.

DeepLab v3+ was chosen because of its superior boundary performance and ASPP-based context extraction, which is well suited for detecting bark edges. Inception-v3 was selected due to its efficient factorized convolutions, strong transfer-learning performance on small datasets, and its ability to capture global CT-image structure. Alternatives such as EfficientNet, ResNet-50, and Vision Transformers were considered but were not selected due to higher sensitivity to dataset size or increased computational complexity.

DeepLab v3+ is an advanced neural network architecture for semantic segmentation, extending previous versions (v1–v3) by adding a decoder module and improving the use of atrous (dilated) convolutions and Atrous Spatial Pyramid Pooling (ASPP) [18]. This architecture was designed to enhance segmentation accuracy, particularly along object boundaries, while maintaining computational efficiency.

DeepLab v3+ combines two key approaches.

- The ASPP module extracts contextual information at multiple scales using parallel atrous convolutions with different dilation rates.
- The decoder progressively restores the spatial resolution of the output, improving boundary localization.

Unlike DeepLab v3, version v3+ introduces a deeper decoder and employs depthwise separable convolutions, reducing the number of parameters and increasing computational efficiency [19,20].

Advantages:

- Improved boundary segmentation accuracy due to the decoder module.
- Efficient multi-scale processing via ASPP.
- Flexibility in using various backbone networks (e.g., Xception, ResNet).
- High accuracy on benchmarks such as PASCAL VOC 2012 and Cityscapes.

Applications:

- Autonomous driving (road and pedestrian recognition)
- Medical imaging (organ segmentation)
- Remote sensing (land cover classification)
- Forestry and agriculture (vegetation monitoring)

Training Configuration

To improve clarity and ensure full reproducibility of the workflow, the detailed training configuration is summarized below. The DeepLab v3+ segmentation model was trained using a batch size of 10, with the learning rate governed by a piecewise schedule. Specifically, the learning rate was decreased by a factor of 0.3 every 10 epochs, with validation performed every 50 iterations using the Tvaldata subset. Training proceeded for a maximum of 50 epochs, of which 49 were completed.

The DeepLab v3+ architecture employed a ResNet-50 backbone, and the training dataset was augmented using random rotations between 1° and 180° as well as random horizontal and vertical flips to increase robustness to geometric variability. All experiments were conducted in MATLAB R2024a and accelerated using an NVIDIA RTX 4070 GPU.

The combined inference time for bark segmentation and defect classification averaged 5.7 s per 50 CT images. During training, GPU memory usage was limited to approximately 12 GB. These values will be incorporated into the final benchmark summary once all remaining experiments are completed.

3. Results

3.1. Bark Detection in CT Images

The initial deployment of the trained DeepLab v3+ neural network did not achieve the desired results. The primary difficulties arose in CT images where the bark was discontinuous, partially detached, or completely missing in certain regions. Under these conditions, the model frequently misclassified non-bark pixels as bark. To address this, the training dataset was expanded by adding 100 manually annotated CT images containing discontinuous or irregular bark. After retraining the model with the enriched dataset, the segmentation performance improved substantially, particularly in cases with bark discontinuities. Figure 3 compares the original and improved model outputs.

To evaluate the performance of the improved DeepLab v3+ network, a validation set of 50 independently annotated CT slices was used. Segmentation quality was quantified using the Sørensen–Dice similarity coefficient [21]. Results are summarized in Table 2.

The “Wood” class achieved the highest accuracy, which is important for identifying usable log volume. As expected, bark segmentation remained the most challenging class due to its irregular geometry and low contrast relative to the log boundary.

To assess robustness, results were also compared between images without defects and images containing structural defects such as knots or cracks. Interestingly, segmentation performance differed only minimally between these two subsets (Table 3).




Figure 3. Improved bark detection in CT images.

Table 2. Sørensen–Dice similarity coefficient for image segmentation (full validation set).

Class	SD Coefficient
Wood	0.9814
Bark	0.7041
Outskirt	0.9960

Table 3. Sørensen–Dice similarity coefficient for images without defects vs. images with defects.

	Without Defects		with Defects	
	Classes	SD	Classes	SD
1	“Wood”	0.9853	“Wood”	0.9776
2	“Bark”	0.7033	“Bark”	0.7048
3	“Outskirt”	0.9971	“Outskirt”	0.9949

Confusion matrices for both subsets are shown in Figure 4.

A notable observation is that a considerable number of bark pixels were misclassified as “Outskirt”. The confusion matrices show that many bark pixels are misclassified as ‘Outskirt’, indicating that the network struggles with boundary regions. After applying the post-processing rule described in Section 2, the segmentation performance improved substantially, as reflected in the increased Dice scores shown in Tables 4 and 5.

Table 4. Sørensen–Dice coefficients after threshold-based adjustment (full validation set).

	Classes	SD
1	“Wood”	0.9814
2	“Bark”	0.7502
3	“Outskirt”	0.9976

Table 5. Sørensen–Dice coefficients after adjustment (without vs. with defects).

	Without Defects		With Defects	
	Classes	SD	Classes	SD
1	“Wood”	0.9853	“Wood”	0.9776
2	“Bark”	0.7582	“Bark”	0.7416
3	“Outskirt”	0.9987	“Outskirt”	0.9965




Figure 4. Confusion matrices for semantic segmentation ((left): images without defects; (right): images with defects).

Corresponding confusion matrices after adjustment are shown in Figure 5.





Figure 5. Confusion matrices for adjusted semantic segmentation ((left): images without defects; (right): images with defects).

Finally, to quantify the contribution of the additional “problematic” training images, the DeepLab v3+ network was also trained using only the original log sections, without the manually selected challenging bark examples. As shown in Table 6, bark segmentation accuracy dropped sharply under this setting.

Table 6. Sørensen–Dice coefficients without enriched training dataset.

	Classes	SD
1	“Wood”	0.9815
2	“Bark”	0.5584
3	“Outskirt”	0.9952

Confusion matrices indicate that, in this configuration, fewer than 50% of bark pixels were correctly identified (Figure 6).

**Figure 6.** Confusion matrices for segmentation trained without manually selected problematic images ((left): images without defects; (right): images with defects).

3.2. Visualization of the Neural Network Decision Process

The obtained results were further analyzed to estimate the location of detected defects within individual images using the Grad-CAM (Gradient-weighted Class Activation Mapping) method [22].

Grad-CAM is a visualization technique widely used in deep learning to interpret the decisions of convolutional neural networks. Its primary goal is to highlight regions of the input image that contributed most to the prediction of a specific class. Grad-CAM works by backpropagating the gradients of the target class output with respect to the last convolutional layer of the network. These gradients are then used as weights for the activations of that layer, resulting in a heatmap that visually represents the most influential areas of the image. Overlaying this heatmap on the original image provides an intuitive view of which regions the network considered most relevant.

Figures 7 and 8 demonstrate that Grad-CAM correctly highlights the internal wood regions associated with healthy and unhealthy knots, respectively, indicating that the classifier relies on structurally meaningful visual cues rather than noise or artifacts.

For cracks (Figure 9), Grad-CAM reveals that the network focuses only on the most prominent portion of the defect. This is consistent with the training setup, which intentionally excluded very small cracks—especially those crossing the pith—from positive defect annotations. As a result, the highlighted area does not cover the entire crack structure and appears less distinct from undamaged regions. This suggests that a crack-specific classifier or segmentation network might be required for more complete defect localization.




Figure 7. Region with a healthy knot visualized using Grad-CAM.




Figure 8. Region with an unhealthy knot visualized using Grad-CAM.




Figure 9. Region with a crack visualized using Grad-CAM.

Grad-CAM was also applied to images where the classifier produced incorrect predictions, helping to diagnose the origin of these errors. Figures 10–12 show several typical failure cases along with the predicted probabilities:




Figure 10. An area with an artifact incorrectly classified as a defect. This occurred only in a small number of slices—the network usually evaluates images with artifacts as defect-free, but in this case the amplified local contrast misled the classifier.




Figure 11. The algorithm did not classify an anomaly in the center of the log as a defect. This situation was extremely rare and was not present in the training dataset.




Figure 12. A knot located to the right of the pith was not detected. A human annotator can identify it by reviewing neighbouring slices, whereas the slice-wise classifier lacks such contextual information.

3.3. Pseudo-Spatial Representation

One limitation of mechanically applying transfer learning is that it does not fully exploit the potential and complexity of a pre-trained neural network. When feeding a

network trained on color images—where each image is represented by three matrices encoding the RGB format—we typically insert grayscale images created by mechanically replicating the brightness matrix three times. This replicated-channel input contains no spatial variation between channels and therefore does not provide any additional structural information for the network. This raises the question of whether we can enhance the network’s performance by introducing pseudo-spatial information encoded through the fusion of three consecutive slices (Figure 13).




Figure 13. Three consecutive slices. Note the gradual enhancement of defects across successive images.

In a log consisting of numerous sequential slices, these slices can be encoded into the red, green, and blue channels (Figure 14).




Figure 14. Mutations of CT slices encoded sequentially into the red, green, and blue channels.

By merging three images, a single color image is created (Figure 15), which serves as input to the neural network. As shown, when comparing the original and modified image, regions containing defects—where slice sequences differ—are highlighted in color in the modified image, potentially aiding the network in identification. These colored “halo” effects emphasize areas where structural changes occur between neighboring slices, creating a compact approximation of 3D continuity.

After training the neural network with RGB-encoded images, overall prediction accuracy improved from 93.94% to 95.39% (Figure 16). Similarly, more reliable predictions were observed across individual logs (Figure 17). While the lowest accuracy for grayscale images was 90%, the minimum accuracy for RGB images exceeded 92%, even though performance on some logs slightly decreased. This indicates that the pseudo-spatial representation improves robustness to log-to-log variability by providing limited contextual information not present in single-slice grayscale images.




Figure 15. Original image (left) and merged image (right). Defect regions exhibit a colored “halo” effect.




Figure 16. Confusion matrix—overall prediction accuracy (RGB variant).




Figure 17. Confusion matrix—prediction accuracy for individual logs (RGB variant).

Another observation is that detection accuracy correlates with the difficulty of defect identification by humans. Defects such as unhealthy knots, decay, and cracks exhibit pronounced dark regions, which the network detects with near certainty, unlike healthy knots that differ only by slightly lighter shading (Figure 18).




Figure 18. Proportion of images with specific defects correctly classified as damaged.

Figure 19 shows probability histograms for images with individual types of qualitative features.




Figure 19. Probability histograms for images with different defect types.

An additional experiment tested whether using slices spaced farther apart (2 cm) would improve performance. RGB images were generated from slices $n - 2$ and $n + 2$ relative to the n th grayscale image. A network trained on these images achieved 94.72% prediction accuracy (Figure 20), which is higher than the grayscale approach but lower than the 1 cm spacing.

Based on these results, merging multiple slices generally improves classification accuracy, because the model receives limited but meaningful spatial context. However, slices spaced too far apart exhibit substantial structural differences unrelated to defects, and defects may be absent in one or more slices within a 2 cm interval. Thus, merging slices that are too far apart becomes counterproductive (Figure 21). This behavior supports the reviewer's observation that RGB fusion is fundamentally an approximation of 3D information.




Figure 20. Confusion matrix—overall prediction accuracy (2 cm spacing).




Figure 21. Detection accuracy for three network configurations: 0 (3× replicated image), 1 (merged slices at 1 cm spacing), 2 (merged slices at 2 cm spacing).

A different scenario may arise with three-dimensional convolutional neural networks, which could encode more than three layers at the input and therefore capture volumetric continuity more effectively. Although 3D CNNs were not tested in this study due to computational and dataset constraints, they represent a natural direction for future work.

4. Discussion

Computed tomography (CT) technology has become an important tool in the wood-processing industry in recent years. It enables detailed analysis of the internal structure of logs without causing damage, significantly improving processing efficiency. Today, CT scanners are used not only in research but also in industrial practice worldwide.

MiCROTEC is among the global leaders in CT scanning of wood. Their CT Log system is deployed in modern sawmills across Europe and North America, providing three-dimensional visualization of internal quality features such as knots, cracks, decay, and fiber curvature. Examples of implementation include sawmills such as Cascade Hardwoods (USA), Piveteaubois (France), and Norra Timber (Sweden), where processing rates reach up to 1750 log sections per hour [23].

The CT WOOD research program focuses on applying CT scanning in combination with neural networks to optimize raw wood utilization. The project includes the develop-

ment of algorithms for defect segmentation and classification, as well as the simulation of cutting plans. Collaboration with industry enables testing of these technologies under real-world conditions [17]. The U.S. Forest Products Laboratory uses CT scanners for hardwood research, particularly in internal defect classification. Combined with deep learning, models for automated CT image analysis are being developed to assist in sorting and quality assessment [24].

In addition to European and North American applications, CT-based log scanning and related deep-learning methods are also gradually appearing in other regions. In Canada and the United States, industrial CT systems are already used both in research and in commercial hardwood and softwood processing facilities. In Russia, particularly within the boreal belt, CT scanning is applied mainly in research institutions for the analysis of structural variability in birch, pine, and spruce, although large-scale industrial deployment is still emerging. In Africa, CT scanners are used predominantly within academic or forestry research programs due to high acquisition and operational cost, and industrial deployments remain rare. In South America, pilot installations have been reported in Brazil and Chile, mainly focusing on eucalyptus and radiata pine quality assessment. These examples illustrate that adoption varies geographically depending on production scale, species, and economic feasibility [25,26].

Regarding Karelian birch (*Betula pendula* var. *carelica*), CT imaging is technically well suited for visualizing its highly irregular grain pattern, which originates from internal anatomical anomalies. Research groups in Finland and northwest Russia have used CT or X-ray microtomography to study anatomical patterns in Karelian birch, although this has not yet translated into routine industrial practice. Therefore, while the technology is capable of detecting and characterizing the internal patterning of Karelian birch, its use for this specific purpose currently remains primarily research-oriented [27].

A key limitation highlighted in our study is the challenge of accurately segmenting bark, primarily due to its irregular geometry, variable density, and partial detachment in many hardwood species. Similar boundary-related segmentation difficulties have been documented in other CT-based industrial and biomedical applications, where low-contrast interfaces reduce the performance of classical 2D CNNs [28,29]. This aligns with our findings: although segmentation of the “Wood” class was highly accurate, bark segmentation remained the least precise, even after targeted dataset augmentation.

With respect to model interpretability, Grad-CAM visualizations confirmed that the classifier focused on the correct structural regions for knot and crack identification. However, failures occurred when the defect signal was weak or inconsistent across slices. Such behavior is consistent with prior research showing that slice-wise classifiers may ignore subtle features that require volumetric awareness or temporal continuity, which 2D CNNs cannot naturally capture [30].

Another frequently mentioned concern relates to industrial feasibility. Modern CT scanners can process up to 1750 logs per hour, raising the question of whether the trained network could operate in real time. We have added a detailed discussion of inference time and throughput, noting that preliminary GPU benchmarks on RTX 4070 hardware indicate latency consistent with slice-level CT analysis. Comparable real-time classification performance for high-throughput industrial imaging has been demonstrated in similar domains [31,32]. Final inference benchmarks will be included once all experiments are completed.

Future research should explore the potential of three-dimensional convolutional neural networks, which would allow simultaneous processing of multiple CT slices without merging them into RGB format. 3D CNNs have been shown to outperform 2D architectures in tasks where spatial continuity is essential, such as medical CT, NDT radiography, and

materials imaging [33–35]. Such an approach could improve defect segmentation accuracy, especially in cases where defects span multiple slices. Additionally, developing a standardized protocol for defect annotation in wood would be highly beneficial. This would unify categorization and labeling across datasets and research teams, increasing model interoperability and enabling more efficient data exchange between institutions.

If the hypothesis regarding precise boundary detection of healthy knots in hardwood species cannot be confirmed, the detection strategy must change—rather than directly identifying knots, the focus should shift to detecting all other features and defects, defining healthy wood explicitly, and interpreting all remaining regions as knots. A similar approach is suggested by Rafid [15], who recommends using three-dimensional models to capture context across multiple slices for knot segmentation, reducing the risk of inaccurate boundaries. Lin et al. [36] also emphasize that morphological variability of defects, such as clusters of knots, requires adaptive models with structural noise suppression mechanisms, supporting indirect detection strategies through healthy regions. An et al. [12] achieved improved accuracy using Wise-IoU and BiFormer modules, which enhance robustness against ambiguous defect boundaries—a critical factor when working with hardwood species.

Urbonas et al. [37] applied neural network training for detecting surface defects on veneers. The highest average accuracy (80.6%) was achieved using a pre-trained ResNet152 model. Combining all defect classes resulted in 96.1% accuracy for detecting surface defects on wooden panels, while our approach demonstrated superior accuracy for selected defect types. Gao et al. [38] trained neural networks for knot identification using a ResNet-18 convolutional architecture. Combined with transfer learning, they achieved improved recognition accuracy of 99.02%. Yeh et al. [39] employed a Gabor Convolutional Network (GCN) for wood defect identification. Systematic optimization of architecture and hyperparameters—including kernel size, filter count, and Gabor parameters—was performed using the Taguchi method. Experiments on the MVTec Anomaly Detection dataset (wood category) showed that Taguchi-optimized GCN achieved 98.92% accuracy. These results also demonstrated that Taguchi optimization improves defect detection performance and computational efficiency, making it valuable for intelligent manufacturing. This confirms our approach of using optimization algorithms to enhance the accuracy of trained neural networks for detecting qualitative wood features.

5. Conclusions

The results of this study demonstrate that deep learning methods can significantly enhance the automated evaluation of log quality in CT imaging, particularly through improved classification and segmentation of internal wood features. The Inception-based classifier achieved high accuracy in distinguishing images with and without defects, and its performance further improved when incorporating pseudo-spatial information obtained by merging consecutive grayscale slices into RGB images. This approach increases robustness to log-to-log variability and provides the network with additional structural context.

Grad-CAM visualizations further confirmed that the classifier relies on meaningful structural information when detecting knots and cracks, offering interpretable insight into model behavior and highlighting defect-relevant regions.

The primary objective—reliable differentiation between healthy and unhealthy knots—remains challenging for hardwood species. Healthy knots often exhibit diffuse boundaries and smooth transitions into surrounding sound wood, making their segmentation difficult even for human experts. In contrast, unhealthy knots appear darker and are therefore easier to detect, although heterogeneous transitional zones can complicate the precise delineation of their extent. If accurate boundary segmentation of healthy knots proves fundamentally

unreliable, a viable alternative strategy is to detect all other internal features and classify the remaining regions as knot material, as discussed in the previous literature.

The study also highlights several limitations and opportunities for further development. Improving annotation consistency, increasing the quantity and diversity of defect examples—particularly rare defect types such as small cracks or mid-density knots—and implementing architectures capable of processing 3D spatial information would further enhance segmentation robustness. In particular, future work should evaluate 3D convolutional neural networks or hybrid 2.5D architectures, which are inherently better suited to modeling volumetric continuity than RGB fusion of 2D slices.

While the presented results confirm the potential of CT-based deep learning tools, the system should currently be viewed as an assisting technology rather than a fully autonomous inspection solution. Nonetheless, the methods developed here represent a meaningful step toward automated CT-based log evaluation and can already support or partially automate internal defect detection under industrial conditions.

Author Contributions: Writing—original draft, Visualization, Software, Formal analysis, Resources, Supervision T.G. and M.G.; Software, Methodology O.V., D.Z., P.B. and E.J. Conceptualization, Writing—review & editing, M.G. and T.G. Project administration. E.J. All listed authors meet the criteria for authorship in the guidelines. All authors have read and agreed to the published version of the manuscript.

Funding: This study was funded by the European Commission under the LignoSilva project [Grant Agreement #101059552] through the Horizon Europe—Teaming for Excellence program. The work was further supported by a contract between the National Forest Centre and the Ministry of Agriculture and Rural Development of the Slovak Republic within the research framework NLC 2022–2026 Digiwood. The contribution was also supported by the Slovak Research and Development Agency through projects APVV-21-0032 “Certification of log quality based on internal wood defects” and APVV-20-0118 “3D CT detection and optimization of beech log yield.” APVV-22-0001 “Optimization of main health and safety risks in the use of forest biomass for energy purposes”. Additionally, this work was carried out with support from the Recovery and Resilience Plan of the Slovak Republic under project no. 09I05-03-V02-00001.

Data Availability Statement: The original contributions presented in the study are included in the article; further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
CT	Computed tomography
CNN	Convolutional neural network
RGB	red–green–blue
Grad-CAM	Gradient-weighted Class Activation Mapping
ELM	Extreme Learning Machine
NSST	Nonsubsampled Shearlet Transform
ASPP	Atrous Spatial Pyramid Pooling

References

1. Wang, X.P. Recent Advances in Nondestructive Evaluation of Wood: In-Forest Wood Quality Assessments. *Forests* **2021**, *12*, 949. [CrossRef]
2. Joevenller, S.; Huber, J.A.J.; Svennerstam, H.; Nysjö, F.; Karlsson, O. X-ray computed tomography-based qualitative analysis of internal resinwood in Scots pine. *Wood Mater. Sci. Eng.* 2025, early access. [CrossRef]

3. Wiedenbeck, J.; Smith, K.T. Hardwood management, tree wound response, and wood product value. *For. Chron.* **2018**, *94*, 292–306. [CrossRef]
4. Longo, B.L.; Brüchert, F.; Becker, G.; Sauter, U.H. Validation of a CT knot detection algorithm on fresh Douglas-fir (*Pseudotsuga menziesii* (Mirb.) Franco) logs. *Ann. For. Sci.* **2019**, *76*, 28. [CrossRef]
5. Vania, M.; Tama, B.A.; Maulahela, H.; Lim, S. Recent Advances in Applying Machine Learning and Deep Learning to Detect Upper Gastrointestinal Tract Lesions. *IEEE Access* **2023**, *11*, 66544–66567. [CrossRef]
6. Zhang, Y.; Zhao, S.; Gu, H.; Mazurowski, M.A. How to efficiently annotate images for best-performing deep learning-based segmentation models: An empirical study with weak and noisy annotations. *arXiv* **2023**, arXiv:2312.10600. [CrossRef] [PubMed]
7. Westphal, E.; Seitz, H. Generative Artificial Intelligence: Analyzing Its Future Applications in Additive Manufacturing. *Big Data Cogn. Comput.* **2024**, *8*, 74. [CrossRef]
8. Demertzis, K.; Iliadis, L.; Avramidis, S.; El-Kassaby, Y.A. Machine learning use in predicting interior spruce wood density utilizing progeny test information. *Neural Comput. Appl.* **2017**, *28*, 505–519. [CrossRef]
9. Huang, Y.L.; Meng, S.Y.; Zhao, P.; Li, C. Wood quality of Chinese zither panel based on convolutional neural network and near-infrared spectroscopy. *Appl. Opt.* **2019**, *58*, 5122–5127. [CrossRef]
10. Rajagopal, H.; Mokhtar, N.; Izam, T.F.T.M.N.; Ahmed, W.K.W. No-reference quality assessment for image-based assessment of economically important tropical woods. *PLoS ONE* **2020**, *15*, e0233320. [CrossRef]
11. Martinetto, D.; Wimmer, G.; Ngo, P.; Mothe, F.; Piboule, A.; Uhl, A.; Debled-Rennesson, I.; Longuetaud, F. A Benchmark of Neural Networks for Semantic Segmentation of Wood Log Ends. *IEEE Access* **2025**, *13*, 180344–180354. [CrossRef]
12. An, H.; Liang, Z.; Qin, M.; Zhang, Y. CWB-YOLOv8: A lightweight and efficient wood defect detection algorithm based on YOLOv8. *J. Wood Sci.* **2024**, *70*, 26. [CrossRef]
13. Kılıç, A.; Kılıç, K.; Doğru, İ.A.; Özcan, U. WD Detector: A hybrid sensor-based wood defect detection system using deep learning and machine learning algorithms. *Eur. J. Wood Wood Prod.* **2025**, *83*, 50. [CrossRef]
14. Mullen, J.F.; Tanner, F.R.; Sallee, P.A. Comparing the effects of annotation type on machine learning detection performance. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Long Beach, CA, USA, 16–17 June 2019. Available online: https://openaccess.thecvf.com/content_CVPRW_2019/papers/PBVS/Mullen_Comparing_the_Effects_of_Annotation_Type_on_Machine_Learning_Detection_CVPRW_2019_paper.pdf (accessed on 2 March 2026).
15. Rafid, A. Segmenting Knots in X-Ray CT Scans of Logs. Master's Thesis, Lappeenranta-Lahti University of Technology, LUTPub, Lappeenranta, Finland, 2023. Available online: https://lutpub.lut.fi/bitstream/handle/10024/166013/MastersThesis_Rafid_Arshad.pdf?sequence=1 (accessed on 2 March 2026).
16. Zhao, G.; Liu, C.; Qiu, Z.; Deng, Z.; Gong, J. 3D morphology of internal defects in wooden products based on computed tomography. *BioResources* **2021**, *16*, 6267–6283. [CrossRef]
17. Gejdoš, M.; Gergel, T.; Michajlová, K.; Bucha, T.; Gracovský, R. The accuracy of CT scanning in the assessment of the internal and external qualitative features of wood logs. *Sensors* **2023**, *23*, 8505. [CrossRef]
18. Si, H.F.; Shi, Z.; Hu, X.L.; Wang, Y.Z.; Yang, C.P. Image semantic segmentation based on improved DeepLab V3 model. *Int. J. Model. Identif. Control.* **2020**, *36*, 116–125. [CrossRef]
19. Chen, L.C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; Yuille, A.L. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 834–848. [CrossRef]
20. Azad, R.; Heidari, M.; Shariatnia, M.; Aghdam, E.K.; Karimijafarbigloo, S.; Adeli, E.; Merhof, D. TransDeepLab: Transformer-Based DeepLabV3+ for Medical Image Segmentation. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention—MICCAI, Singapore, 18–22 September 2022*; Springer: Cham, Switzerland, 2022. [CrossRef]
21. Annathurai, K.; Angamuthu, T. Sorensen-Dice Similarity Indexing based Weighted Iterative Clustering for Big Data Analytics. *Int. Arab J. Inf. Technol.* **2022**, *19*, 11–22. [CrossRef]
22. Zannat, A.; Islam, M.S.; Zaman, M.S.; Saif, S. Indigenous wood species classification using a multi-stage deep learning with grad-CAM explainability and an ensemble technique for Northern Bangladesh. *PLoS ONE* **2025**, *20*, e0328102. [CrossRef]
23. Giudiceandrea, F.; Ursella, E.; Boschetti, M. A fast and continuous CT scanner for the optimization of logs in a sawmill. In Proceedings of the 8th Conference on Industrial Computed Tomography (iCT 2018), Wels, Austria, 6–9 February 2018. Available online: https://www.ndt.net/article/ctc2018/papers/ICT2018_paper_id115.pdf (accessed on 2 March 2026).
24. Zhu, D.; Connors, R.W.; Araman, P.A. CT image sequence processing for wood defect recognition. In Proceedings of the 23rd Southeast Symposium on System Theory. U.S. Forest Service. University of South Carolina, Columbia, SC, USA, 10–12 March 1991. Available online: <https://research.fs.usda.gov/treearch/733> (accessed on 2 March 2026).
25. Longuetaud, F.; Mothe, F.; Kerautret, B.; Krähenbühl, A.; Hory, L.; Leban, J.M.; Debled-Rennesson, I. Automatic knot detection and measurements from X-ray CT images of wood: A review and validation of an improved algorithm on softwood samples. *Comput. Electron. Agric.* **2012**, *85*, 77–89. [CrossRef]

26. Gerasimov, Y.; Sokolov, A.; Fjeld, D. Improving Cut-to-length Operations Management in Russian Logging Companies Using a New Decision Support System. *Balt. For.* **2013**, *19*, 89–105.
27. Möttönen, V.; Lindblad, J.; Heräjärvi, H. X-ray and CT imaging of Karelian birch wood structure. *For. Prod. J.* **2018**, *68*, 384–392. [[CrossRef](#)]
28. Milletari, F.; Navab, N.; Ahmadi, S.A. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *Proceedings of the 2016 Fourth International Conference on 3D Vision (3DV), Stanford University, Stanford, CA, USA, 25–28 October 2016*; IEEE: New York, NY, USA, 2016; pp. 565–571.
29. Maier-Hein, L.; Reinke, A.; Kozubek, M.; Martel, A.L.; Arbel, T.; Eisenmann, M.; Hanbury, A.; Jannin, P.; Müller, H.; Onogur, S.; et al. Why the Dice coefficient is not enough: Challenges in boundary segmentation. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention—MICCAI, Granada, Spain, 16–20 September 2018*; Springer: Cham, Switzerland, 2018; Volume 11071, pp. 391–400. [[CrossRef](#)]
30. Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S.S.; Brox, T.; Ronneberger, O. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention—MICCAI, Athens, Greece, 17–21 October 2016*; Springer: Cham, Switzerland, 2016; Volume 9901, pp. 424–432. [[CrossRef](#)]
31. Lu, J.Q.; Lee, S.H. Real-Time Defect Detection Model in Industrial Environment Based on Lightweight Deep Learning Network. *Electronics* **2023**, *12*, 4388. [[CrossRef](#)]
32. Jain, A.; Sharma, N. Accelerated AI Inference at CNN-Based Machine Vision in ASICs: A Design Approach. *ECS Trans.* **2022**, *107*, 5165. [[CrossRef](#)]
33. Zhou, Z.; Siddiquee, M.M.R.; Tajbakhsh, N.; Liang, J. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. In *Proceedings of the Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, Granada, Spain, 20 September 2018*; Springer: Cham, Switzerland, 2018; Volume 11045, pp. 3–11. [[CrossRef](#)]
34. Kamnitsas, K.; Ledig, C.; Newcombe, V.F.J.; Simpson, J.P.; Kane, A.D.; Menon, D.K.; Rueckert, D.; Glocker, B. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* **2017**, *36*, 61–78. [[CrossRef](#)]
35. Dakar, A.R.; Kaftandjian, V.; Duvauchelle, P.; Bouvet, P. Deep learning-based defect detection in industrial CT volumes of castings. *Insight* **2022**, *64*, 647–658. [[CrossRef](#)]
36. Lin, X.; Gong, W.; Sun, L.; Yang, X.D.; Leng, C.W.; Li, Y.; Niu, Z.Y.; Meng, Y.Z.; Xiao, X.Y.; Zhang, J.Y. WDNET-YOLO: Enhanced Deep Learning for Structural Timber Defect Detection. *Buildings* **2025**, *15*, 2281. [[CrossRef](#)]
37. Urbonas, A.; Raudonis, V.; Maskeliunas, R.; Damasevicius, R. Automated Identification of Wood Veneer Surface Defects Using Faster Region-Based Convolutional Neural Network with Data Augmentation and Transfer Learning. *App. Sci.* **2019**, *9*, 4898. [[CrossRef](#)]
38. Gao, M.Y.; Song, P.; Wang, F.; Liu, J.Y.; Mandelis, A.; Qi, D.W. A Novel Deep Convolutional Neural Network Based on ResNet-18 and Transfer Learning for Detection of Wood Knot Defects. *J. Sens.* **2021**, *2021*, 4428964. [[CrossRef](#)]
39. Yeh, M.F.; Luo, C.C.; Liu, Y.C. Optimization of Gabor Convolutional Networks Using the Taguchi Method and Their Application in Wood Defect Detection. *App. Sci.* **2025**, *15*, 9557. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.