International Journal of Engineering Research & Science (IJOER)
ISSN: [2395-6992]
[Vol-12, Issue-8, August - 2026]
Website: www.ijoer.com
Journal DOI: 10.25125/engineering-journal

Laboratory Validation of a Low-Cost Embedded Computer Vision System for Automated Defect Detection in Beech Sawn Timber

Sorin Eugen Popa[1]; Roxana Margareta Grigore[2*]

Vasile Alecsandri University of Bacău, Faculty of Engineering, Department of Power Engineering and Computer Science, Bacău, Romania
*Corresponding Author

Received: 29 July 2026/ Revised: 03 August 2026/ Accepted: 10 August 2026/ Published: 15-08-2026

***Abstract*— *The digital transformation of the wood-processing industry increasingly relies on computer vision, artificial intelligence (AI), and embedded sensing technologies to automate timber quality assessment. This study presents the laboratory validation (TRL4) of a low-cost embedded computer-vision demonstrator for automated surface-defect detection in beech (*Fagus sylvatica*) sawn timber, developed as the first operational module of the SMARTWOOD-AI technology-transfer platform. The system integrates a Raspberry Pi 4 single-board computer with a Sony IMX500 intelligent camera and employs an interpretable computer-vision pipeline based on 21 handcrafted colour (HSV), texture (Gray-Level Co-occurrence Matrix and Local Binary Patterns), and edge-density (Canny) features classified using a Random Forest algorithm. A dataset comprising 24 beech boards (192 labelled image regions) was evaluated using a board-level train/validation partitioning strategy to prevent data leakage. The classifier achieved an average accuracy of 88.2% (±7.1%) during five-fold cross-validation and 82.5% accuracy on an independent validation set, with high sensitivity for defect detection (recall = 0.93, F1-score = 0.88). The results demonstrate the technical feasibility of the proposed embedded inspection architecture and establish a reproducible experimental baseline for future integration of deep-learning models, digital twins, and intelligent cutting optimisation within the SMARTWOOD-AI platform.*

***Keywords*— wood defect detection; embedded computer vision; on-sensor AI; Random Forest; Technology Readiness Level (TRL4); beech sawn timber; digital twin; edge artificial intelligence; Industry 4.0; smart sawmill.**

I. INTRODUCTION

The digital transformation of the wood-processing industry is reshaping the way timber quality is assessed and production decisions are made. Within the Industry 4.0 and Industry 5.0 paradigms, computer vision, artificial intelligence (AI), edge computing and digital-twin technologies are increasingly integrated to automate inspection processes, improve production efficiency and maximise raw-material utilisation. Recent bibliometric evidence based on more than one thousand scientific publications demonstrates a rapid acceleration of research on AI-assisted wood inspection, with surface quality assessment representing one of the most active application domains [1]. As a consequence, automated defect detection has become a key enabling technology for intelligent sawmills, where reliable visual information can directly support optimisation of cutting strategies and resource utilisation.

Despite these advances, quality grading in many small and medium-sized wood-processing enterprises still relies predominantly on manual visual inspection. Although experienced operators can identify most visible defects, the process remains subjective, operator-dependent and sensitive to fatigue, illumination conditions and the inherent variability of natural wood. These limitations often result in inconsistent grading decisions and inefficient material utilisation [2, 3]. Since natural defects may reduce industrial timber utilisation to only 50–70% [4], the development of reliable automated inspection systems remains both an industrial and an economic priority.

Research on automated wood inspection has evolved considerably during the last two decades. Early approaches relied on hand-crafted image descriptors combined with conventional machine-learning algorithms. Texture descriptors derived from the Gray-Level Co-occurrence Matrix (GLCM), Local Binary Patterns (LBP) and other manually engineered features proved effective for detecting knots and cracks, while ensemble classifiers such as Random Forest offered competitive performance together with computational efficiency and model interpretability [5-9]. Although these approaches have gradually been superseded by deep-learning methods, they remain particularly attractive when only limited annotated datasets are available and when transparent decision-making is required.

During recent years, deep convolutional neural networks (CNNs) and one-stage object detectors of the YOLO family have become the dominant paradigm for wood defect detection [4, 10-15]. Numerous studies have reported excellent detection accuracy through specialised architectures designed for small defects, complex textures and multi-scale feature extraction [4, 10-25]. However, these models generally rely on thousands of annotated images, extensive computational resources and offline training procedures, making them more suitable for mature technological developments than for early-stage laboratory validation. In parallel, increasing attention has been devoted to edge AI and embedded vision, where lightweight models are deployed on low-cost hardware platforms such as Raspberry Pi devices [26, 27]. More recently, intelligent vision sensors such as the Sony IMX500 have introduced the concept of on-sensor AI inference, enabling neural-network execution directly within the image sensor and significantly reducing data-transfer latency and computational load [28].

Although these technological advances have substantially improved automated defect detection, several important research gaps remain. Bibliometric analyses identify real-time industrial integration, edge deployment of AI algorithms and the incorporation of computer vision into digital-twin-based decision systems as emerging research directions that remain insufficiently explored [1]. Furthermore, most published studies focus primarily on maximising classification accuracy using large public datasets, whereas comparatively little attention has been paid to the experimental validation of low-cost embedded inspection systems intended for the early stages of technology maturation. Equally important, many studies perform dataset splitting at image level rather than at board level, introducing data leakage that may artificially inflate reported performance.

Against this background, the present study addresses these gaps by reporting the laboratory validation (TRL4) of a fully embedded computer-vision demonstrator developed within the SMARTWOOD-AI platform. Unlike previous studies focused primarily on algorithmic performance, this work emphasises the technological validation of an integrated inspection chain combining a Raspberry Pi 4 single-board computer with a Sony IMX500 intelligent camera operating under controlled laboratory conditions. The proposed solution deliberately employs an interpretable classical machine-learning pipeline based on twenty-one colour, texture and edge descriptors and a Random Forest classifier, which is particularly suitable for the limited experimental dataset available at this stage of technological maturity. In addition, the study adopts a rigorous board-level data-partitioning protocol that prevents data leakage and provides a realistic evaluation of classifier performance. The obtained results establish the experimental baseline for the subsequent integration of deep-learning models, digital twins and multi-objective cutting optimisation within the SMARTWOOD-AI platform.

II. MATERIAL AND METHODS

2.1 Experimental Platform

Experimental validation of the demonstrator was conducted in the laboratory of the EMI Research Centre, Faculty of Engineering, "Vasile Alecsandri" University of Bacău, using a dedicated experimental platform specifically designed to ensure reproducible image acquisition under controlled environmental conditions. The objective of the experimental setup was to minimise the influence of external factors, such as illumination variability, geometric distortions and background interference, thereby enabling the assessment of the proposed computer-vision pipeline under repeatable laboratory conditions consistent with Technology Readiness Level 4.

The hardware platform consisted of a Raspberry Pi 4 Model B single-board computer, responsible for image acquisition, system control and execution of the computer-vision algorithms, and a Sony IMX500 intelligent vision camera integrating on-sensor AI capabilities within a CMOS image sensor. The camera was configured to acquire images at a resolution of 1920 × 1080 pixels with a fixed exposure time of 15 ms and a frame rate of 10 fps. Although the present study employs a classical machine-learning pipeline executed on the Raspberry Pi, the selection of the IMX500 provides a hardware architecture compatible with future migration towards embedded deep-learning inference without modifications to the acquisition subsystem.

To ensure stable imaging conditions, the experimental platform incorporated a controlled illumination system composed of two symmetrically positioned LED light sources with a colour temperature of 4000–5000 K. The symmetric arrangement minimised shadows and reduced illumination non-uniformity across the inspected timber surface, while the selected lighting conditions are representative of those commonly adopted in industrial machine-vision applications. The radiometric characteristics of the illumination system can be evaluated using the low-cost thermal-sensing methodology previously developed by the authors [29]. A matte black background was employed to maximise contrast between the timber specimen and the surrounding scene, facilitating reliable image segmentation, while a dimensional reference marker positioned within the field of view enabled geometric calibration and dimensional verification of the acquired images.

The Sony IMX500 camera was mounted perpendicular to the timber surface at a fixed working distance of 300 mm, and its position remained unchanged throughout the entire experimental campaign. Maintaining constant acquisition geometry minimised perspective distortions and scale variations between successive images, ensuring that the extracted colour, texture and edge descriptors reflected the intrinsic characteristics of the wood surface rather than variations introduced by the imaging process.

(a) **(b)** **(c)**

FIGURE 1: Experimental image-acquisition stand

(a) Overall view of the experimental stand showing the timber specimen positioned on the matte black background, with the Sony IMX500 camera mounted perpendicular to the surface and the two symmetrically positioned LED light sources; (b) Raspberry Pi 4 Model B single-board computer showing the hardware connections; (c) Sony IMX500 intelligent vision camera mounted on the support structure)

2.2 Experimental Material

The experimental dataset consisted of representative beech (*Fagus sylvatica*) sawn timber samples collected directly from the production flow of the industrial partner participating in the SMARTWOOD-AI project. Using industrial material rather than laboratory-prepared specimens ensured that the analysed surfaces reflected the natural variability of wood encountered in practical manufacturing environments, including differences in colour, grain orientation, machining traces and naturally occurring defects.

A total of 24 sawn boards were included in the experimental campaign. The boards had nominal dimensions of approximately 500 mm × 120 mm × 25 mm (length × width × thickness). Moisture content was measured using a resistance-type moisture meter and was found to be within the range of 8–12%, representative of air-dried industrial timber. Each board was imaged on both faces under identical acquisition conditions. To increase the number of learning samples while preserving the physical independence of the observations, each image was uniformly divided into four non-overlapping regions of equal size, resulting in a dataset of 192 image regions used for classifier development and validation.

The dataset included both defect-free surfaces and surfaces exhibiting the most common defects encountered in industrial beech processing, including knots, longitudinal cracks, discolorations, residual bark and combinations of multiple defect types. Only two of the 24 boards were entirely free of natural defects, reflecting the prevalence of imperfections in industrial timber. Consequently, the experimental dataset captured the diversity of visual characteristics expected during industrial operation while remaining representative of an early-stage technology-validation campaign. Although relatively limited in size, the dataset was considered sufficient for validating the functionality of the proposed demonstrator at Technology Readiness Level 4 (TRL4) and for establishing a reliable experimental baseline for subsequent developments.

2.3 Image Acquisition and Labelling Protocol

All images were acquired using the experimental platform described in the previous section under identical laboratory conditions. Camera position, acquisition geometry, illumination intensity and exposure parameters remained constant throughout the entire experimental campaign in order to minimise the influence of external factors on the extracted image descriptors.

Several image-acquisition strategies were evaluated during the preliminary design phase of the demonstrator. The adopted protocol consisted of photographing each face of every board at full resolution, followed by uniform segmentation into four equally sized image regions. This strategy provided three important advantages. First, it increased the number of available training samples without requiring additional physical specimens. Second, it standardised the spatial dimensions of all analysed images, ensuring consistent feature extraction. Third, it reduced the influence of local defect distribution by allowing individual regions of the same board to be analysed independently.

Each segmented image was manually inspected and annotated using a standardised English nomenclature comprising the classes nodefect, defect_knot, defect_crack, defect_discoloration, defect_bark and defect_mixed. Although the annotation protocol preserved the specific defect category associated with each image, the present study addressed the more fundamental problem of binary defect detection, in which all defect categories were merged into a single defect class while maintaining nodefect as the reference class. This simplification was intentionally adopted because the primary objective of the present work was the experimental validation of the integrated computer-vision demonstrator rather than the development of a multiclass defect-recognition system.

To ensure annotation consistency, all labels were reviewed after the initial dataset construction. During this quality-control procedure, several discrepancies between the preliminary visual assessment and the actual surface condition were identified and corrected through manual reinspection of the corresponding timber samples. This verification step improved dataset reliability and reduced the risk of introducing annotation errors into the machine-learning process.

2.4 Image Pre-processing

Prior to feature extraction, all acquired images underwent a standardised pre-processing pipeline designed to minimise the influence of acquisition-related variability while preserving the visual characteristics relevant for defect detection. The primary objective of the pre-processing stage was to ensure that the subsequent colour, texture and edge descriptors reflected the intrinsic properties of the timber surface rather than variations introduced by the imaging process.

The pre-processing workflow consisted of four consecutive steps. First, the contour of the timber specimen was automatically identified using Otsu's thresholding method, which separated the foreground from the matte black background without requiring manual intervention. Second, small artefacts generated during image acquisition or resulting from mechanical processing were removed by applying morphological opening operations, thereby reducing the influence of isolated noise on the extracted descriptors. Third, the detected contour was used to automatically crop the region of interest and eliminate the surrounding background, ensuring that feature extraction was performed exclusively on the timber surface. Finally, all cropped images were normalised to a common spatial resolution to guarantee consistent feature computation across the entire dataset.

The adopted pre-processing strategy ensured that all image regions were analysed under comparable conditions despite the natural variability of timber geometry. By reducing background interference, geometric inconsistencies and acquisition noise, the pre-processing stage increased the robustness and repeatability of the extracted colour, texture and edge features, thereby improving the reliability of the subsequent machine-learning classification.

2.5 Feature Extraction

Following image pre-processing, each image region was represented by a numerical feature vector describing complementary visual characteristics of the timber surface. The objective of the feature-extraction stage was to transform the raw image into a compact and discriminative representation suitable for machine-learning classification while maintaining computational efficiency and model interpretability, both essential requirements for an embedded TRL4 demonstrator.

A total of 21 hand-crafted features were extracted and grouped into four complementary descriptor families:

- **Colour features** (6 features), represented by the mean and standard deviation of the Hue (H), Saturation (S) and Value (V) channels of the HSV colour space. Colour space conversion from RGB to HSV was performed using the standard OpenCV implementation (cv2.COLOR_BGR2HSV). Compared with the RGB representation, the HSV space provides improved separation between chromatic information and illumination intensity, making it more suitable for analysing natural variations in wood colour.

- **Texture features derived from the Gray-Level Co-occurrence Matrix (GLCM)** (5 features), including contrast, homogeneity, correlation, energy and dissimilarity. GLCM features were computed on grayscale images using a pixel distance of 1 and averaged over four directions (0°, 45°, 90°, 135°), with 256 quantization levels. These second-order statistical descriptors characterise the spatial relationships between neighbouring pixels and have been widely used for discriminating natural wood defects such as knots, cracks and discolorations [5-9].

- **Local Binary Pattern (LBP) descriptors** (6 features), computed using uniform patterns with a radius of 1 and 8 neighbours. Statistical measures derived from the LBP histograms (mean, variance, skewness, kurtosis, entropy, and energy) were extracted as features. LBP descriptors provide an efficient representation of local micro-texture and are particularly robust to moderate illumination variations, making them well suited for characterising the heterogeneous surface structure of sawn timber.

- **Structural information** (4 features), represented by the edge density and three additional statistical descriptors (mean, variance, and number of edge pixels) obtained using the Canny edge detector. The Canny parameters were set to threshold values of 50 and 150 for the lower and upper hysteresis thresholds, respectively. The edge-density feature quantifies the amount of local structural discontinuities associated with cracks, fibre disruptions and defect boundaries.

The combination of these descriptor families enabled the simultaneous representation of chromatic, textural and structural information, providing complementary evidence for distinguishing defective from defect-free timber surfaces. This hybrid feature representation was intentionally selected because it offers a favourable compromise between computational complexity, interpretability and classification performance for relatively small experimental datasets, where feature-based machine-learning methods remain a practical alternative to data-intensive deep-learning approaches.

To further analyse the contribution of each descriptor family, the relative importance of every feature was subsequently estimated using the intrinsic feature-importance mechanism of the Random Forest classifier. This analysis not only supported the interpretation of the classification results but also identified the most informative visual characteristics for future optimisation of the feature set and for comparison with subsequent deep-learning models.

2.6 Classification Model

Image classification was performed using the Random Forest algorithm, an ensemble learning method that combines multiple decision trees through majority voting. Random Forest was selected because it provides a favourable balance between classification performance, robustness and computational efficiency, making it particularly suitable for relatively small experimental datasets composed of heterogeneous colour, texture and structural descriptors. In addition, the algorithm requires only limited parameter tuning, exhibits good generalisation capability and is less susceptible to overfitting than many individual decision-tree models.

The selection of Random Forest was also motivated by the objectives of the present study. Rather than maximising classification accuracy, the primary goal was to validate the functionality of the complete embedded computer-vision demonstrator under Technology Readiness Level 4 conditions using a representative, but relatively limited, experimental dataset. Under these circumstances, feature-based machine-learning methods provide a practical and computationally efficient solution while maintaining full compatibility with low-cost embedded hardware platforms.

Although deep-learning approaches based on Convolutional Neural Networks (CNNs) currently represent the state of the art for automated wood-defect detection [10-15, 18-23], their successful application generally requires substantially larger annotated datasets together with significantly greater computational resources for model training and optimisation. Consequently, the adoption of a Random Forest classifier should be regarded as a deliberate methodological choice aligned with the objectives and maturity level of the proposed demonstrator rather than as a limitation of the developed system.

An additional advantage of the selected classifier is its intrinsic capability to estimate the relative importance of each input feature. This feature-importance analysis provides valuable insight into the contribution of colour, texture and edge descriptors to the classification process, improving model interpretability and supporting the optimisation of future computer-vision models. Such interpretability is particularly valuable in industrial decision-support applications, where transparent and explainable classification results facilitate system validation and user confidence.

The Random Forest classifier was implemented in scikit-learn (version 1.9.0) with 200 trees (n_estimators = 200), the Gini splitting criterion, and no maximum depth constraint (max_depth = None). The number of trees was selected based on preliminary experiments indicating convergence of the out-of-bag error at approximately 150–200 trees. A fixed random seed (random_state = 42) was used to ensure reproducibility. To mitigate the pronounced class imbalance of the training set, class weights were set to "balanced" (class_weight = "balanced"), which inversely weights each class in proportion to its frequency; the remaining hyperparameters were kept at their scikit-learn default values.

2.7 Experimental Validation Protocol

To obtain an unbiased estimate of classifier performance, the experimental dataset was partitioned at the physical-board level rather than at the image level. Consequently, all image regions originating from the same timber board were assigned exclusively to either the training or the validation subset, preventing information leakage between the two datasets. This strategy ensured that the validation process evaluated the classifier on previously unseen physical specimens rather than on visually similar regions extracted from the same board, thereby providing a more realistic assessment of its generalisation capability.

Classifier development and performance evaluation were carried out using five-fold cross-validation on the training dataset, followed by an independent evaluation on the validation set. Cross-validation was employed to assess the stability of the model and to reduce the influence of the relatively limited dataset size, whereas the independent validation set provided an unbiased estimate of the expected performance under practical operating conditions.

Performance was quantified using the standard metrics adopted for binary classification, namely accuracy, precision, recall, F1-score, and the confusion matrix. These complementary indicators allow both the overall classification performance and the class-specific behaviour of the model to be evaluated, providing a comprehensive assessment of its suitability for automated defect detection.

Confidence intervals for the validation accuracy were computed using the Wilson score method, which is appropriate for binomial proportions estimated on small samples.

2.8 Demonstrator Software Architecture

The developed demonstrator integrates the complete computer-vision workflow into a modular embedded architecture comprising five functional components: (i) image acquisition, (ii) image preprocessing and normalisation, (iii) feature extraction, (iv) Random Forest classification, and (v) result visualisation and data storage.

The overall processing workflow, illustrated in Figure 2, begins with image acquisition using the Sony IMX500 intelligent vision camera. The acquired images undergo automatic identification of the timber surface, followed by background removal, geometric normalisation and segmentation into individual analysis regions. For each region, colour, texture and structural descriptors are extracted and assembled into a numerical feature vector, which is subsequently classified by the Random Forest model as either defect or nodefect. Finally, the classification results are displayed and stored for further analysis and integration with the higher-level modules of the SMARTWOOD-AI platform.

The complete software pipeline was implemented in Python and executed on the Raspberry Pi 4 embedded platform. The successful integration of image acquisition, pre-processing, feature extraction and classification within a low-cost embedded system demonstrates the feasibility of performing automated timber inspection without relying on external computing infrastructure. This experimental validation establishes the technological basis for the future integration of deep-learning models, digital-twin functionality and intelligent cutting optimisation within the SMARTWOOD-AI platform.

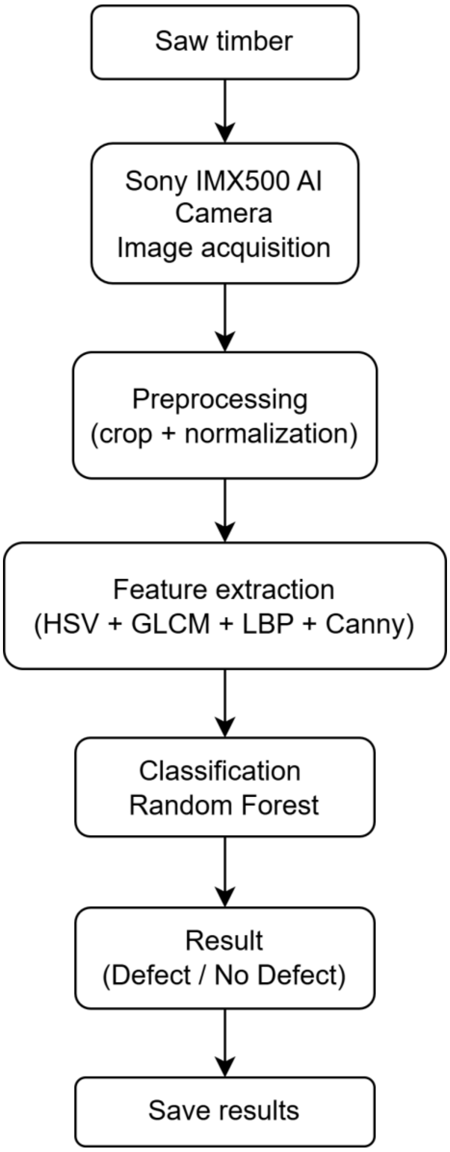

FIGURE 2: Functional processing flow of the experimental demonstrator

III. RESULTS AND DISCUSSION

3.1 Experimental Configuration

The proposed computer-vision demonstrator was experimentally validated using the dataset described in Section 2 under controlled laboratory conditions. The complete dataset comprised 192 image regions extracted from 24 beech (*Fagus sylvatica*) boards, with data partitioning performed at the physical-board level to prevent information leakage between the training and validation subsets.

The Random Forest classifier was trained using the extracted colour, texture and structural descriptors. Model robustness was first assessed through five-fold cross-validation on the training dataset and subsequently verified using an independent validation set composed of previously unseen timber boards. This evaluation strategy provided both an estimate of model stability during training and an unbiased assessment of its generalisation capability.

Table 1 summarises the distribution of defect and defect-free samples used for training and independent validation. As expected for industrial timber, the experimental dataset exhibits a pronounced class imbalance, with defective surfaces representing the majority of the analysed samples. This class distribution is representative of the material selected for the present laboratory validation and constitutes an additional challenge for automated classification.

TABLE 1
COMPOSITION OF THE TRAINING AND VALIDATION SETS

Set	Total images	defect
Training	152	138
Validation	40	28

3.2 Classifier Performance

The classification performance of the proposed demonstrator was evaluated using the standard metrics for binary classification, namely accuracy, precision, recall and F1-score. The Random Forest classifier achieved an average accuracy of 88.2% (±7.1%) during five-fold cross-validation on the training dataset and 82.5% accuracy on the independent validation set (33 of 40 image regions correctly classified; 95% Wilson confidence interval: 68.1–91.3%). The relatively wide confidence interval reflects the limited size of the validation dataset and is consistent with the exploratory nature of this laboratory-scale TRL4 validation. The cross-validation standard deviation of ±7.1% indicates moderate variability across different training/validation splits, which is expected given the limited dataset size and the class imbalance. These results indicate that the proposed embedded computer-vision system is capable of reliably distinguishing defective from defect-free timber surfaces under controlled laboratory conditions.

The lower performance observed on the independent validation set compared with cross-validation is expected because the validation images originate from previously unseen physical boards and therefore provide a more rigorous assessment of model generalisation. Considering the relatively limited size of the experimental dataset and the inherent variability of natural wood, the obtained performance confirms the suitability of the proposed methodology for laboratory validation.

The class-wise evaluation metrics are presented in Table 2. The classifier achieved higher performance for the defect class than for the nodefect class, with a defect-class recall of 0.93, indicating that the majority of defective surfaces were correctly identified. This behaviour is particularly relevant for industrial inspection applications, where minimising missed defects is generally more important than eliminating all false-positive detections.

TABLE 2
CLASS-WISE PERFORMANCE OF THE RANDOM FOREST CLASSIFIER ON THE VALIDATION SET

Class	Precision	Recall
nodefect (n = 12)	0.78	0.58
defect (n = 28)	0.84	0.93

3.3 Confusion Matrix Analysis

The confusion matrix obtained for the independent validation set provides a detailed overview of the classifier performance and the distribution of correctly and incorrectly classified samples (Table 3 and Figure 3). Overall, the classifier correctly identified 33 of the 40 analysed image regions, corresponding to an overall validation accuracy of 82.5%.

TABLE 3
CONFUSION MATRIX OF THE RANDOM FOREST CLASSIFIER

Actual	Nodefect	Defect
Nodefect	7	5
Defect	2	26

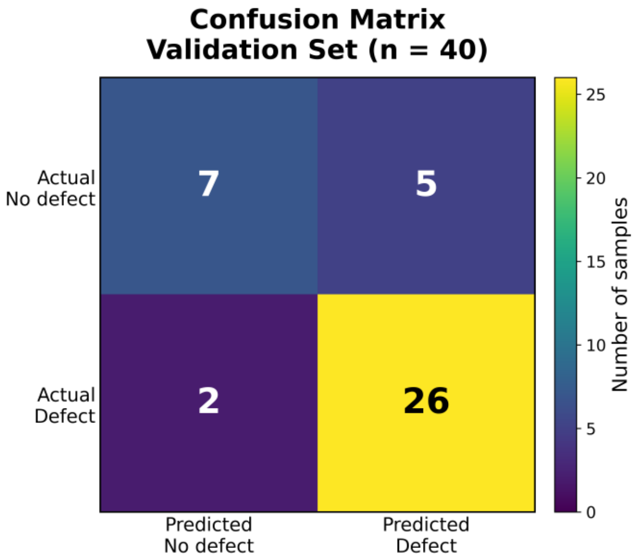

FIGURE 3: Confusion matrix of the Random Forest classifier.

The classifier demonstrated a strong ability to identify defective timber surfaces, correctly recognising 26 of the 28 defect samples, corresponding to a recall of 0.93 for the defect class. Only two defective samples were incorrectly classified as defect-free, indicating a relatively low false-negative rate.

Performance for the nodefect class was comparatively lower, with 7 of the 12 defect-free samples correctly identified and 5 samples incorrectly classified as defective. Consequently, the classifier exhibits a tendency to favour high defect sensitivity at the expense of a larger number of false-positive detections. Such behaviour is consistent with the design objectives of an early-stage inspection system, where overlooking a defective timber piece is generally considered more critical than submitting a sound piece for additional inspection.

Visual examination of the misclassified samples indicates that most classification errors occurred in image regions characterised by subtle defects, pronounced natural texture variability or machining artefacts that locally altered the surface appearance. These observations suggest that the extracted texture descriptors successfully capture structural irregularities associated with defects but may also respond to non-defect surface features presenting similar visual characteristics. A more detailed analysis of the underlying error mechanisms is presented in Section 3.5.

3.4 Feature Importance Analysis

One of the principal advantages of the Random Forest algorithm is its ability to estimate the relative contribution of each input feature to the classification process. The feature-importance analysis, presented in Figure 4, provides insight into the visual characteristics that most strongly influenced the classifier decisions.

The obtained results indicate that texture descriptors contributed most significantly to the classification performance. Features derived from the Gray-Level Co-occurrence Matrix (GLCM) and Local Binary Patterns (LBP) consistently exhibited the highest importance values, demonstrating that local texture information represents the most discriminative characteristic for distinguishing defective from defect-free beech timber surfaces.

The HSV colour descriptors provided a complementary contribution by capturing chromatic variations associated with defects such as discolorations and bark inclusions. Although their overall importance was lower than that of the texture descriptors, colour information contributed to improving the discrimination between visually similar surface regions. The edge-density feature, computed using the Canny detector, showed the smallest contribution, indicating that structural discontinuities alone are insufficient for reliable defect identification but provide useful complementary information when combined with colour and texture descriptors.

Overall, the feature-importance analysis confirms that the adopted hybrid representation successfully captures complementary visual information describing timber surfaces. The predominance of texture descriptors is consistent with the visual characteristics of natural wood defects and supports the selection of GLCM and LBP features as the principal components of the proposed computer-vision pipeline.

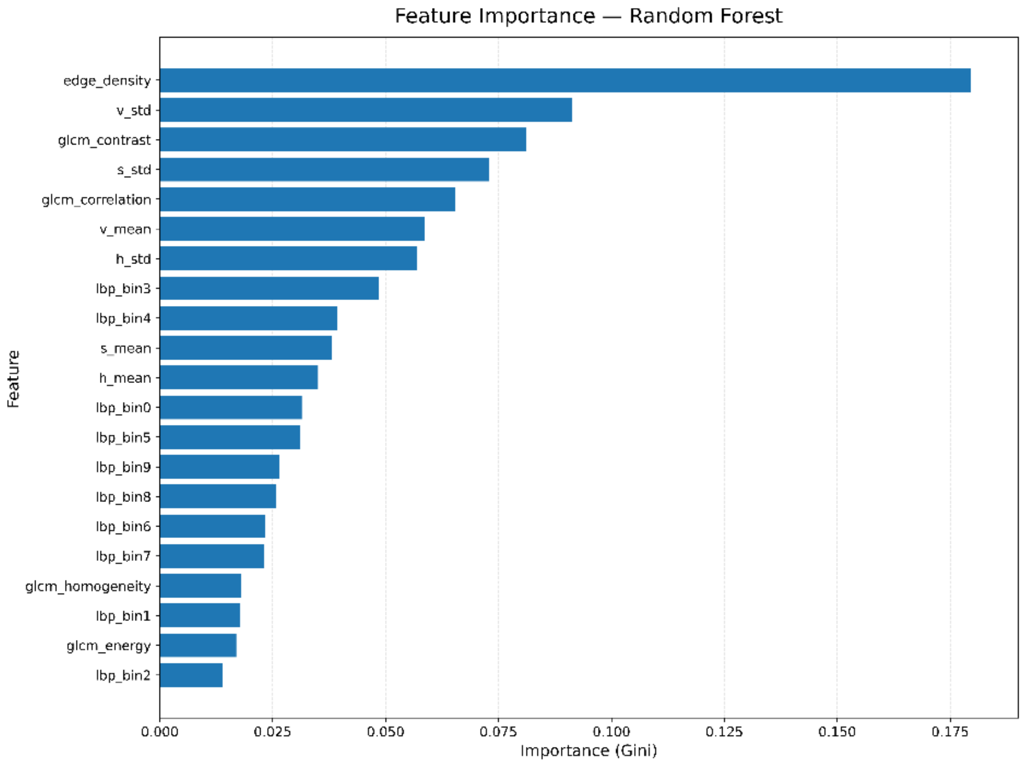

FIGURE 4: Relative importance of the features used by the Random Forest classifier

3.5 Discussion

The present study demonstrates the technical feasibility of a low-cost embedded computer-vision system for the automated identification of defects in beech sawn timber under controlled laboratory conditions. Rather than aiming to maximise classification accuracy, the primary objective was to validate the integration of image acquisition, pre-processing, feature extraction and automatic classification into a single functional demonstrator corresponding to Technology Readiness Level 4 (TRL4). From this perspective, the successful operation of the complete processing chain confirms that the proposed architecture constitutes a reliable technological foundation for the subsequent industrial development of the SMARTWOOD-AI platform.

The objective of the present work was therefore fundamentally different from that of algorithm benchmarking studies. Instead of maximising predictive performance on large public datasets, the emphasis was placed on demonstrating the feasibility, robustness and reproducibility of an integrated embedded inspection architecture suitable for subsequent industrial maturation.

The Random Forest classifier achieved an accuracy of 82.5% on the independent validation set and 88.2% (±7.1%) during five-fold cross-validation. Although these values are lower than those typically reported for recent deep-learning architectures trained on large public datasets [10-15, 18-23], the comparison should be interpreted with caution because the experimental conditions differ substantially. Most state-of-the-art CNN- and YOLO-based studies rely on several thousand annotated images, extensive data augmentation and high-performance computational resources, whereas the present work intentionally focuses on a limited experimental dataset representative of an early technology-development stage. Consequently, the reported performance should not be regarded as the maximum achievable accuracy of the proposed methodology, but rather as evidence that an embedded inspection chain based on classical machine-learning techniques can already provide reliable defect discrimination while maintaining computational simplicity, model interpretability and low hardware requirements.

To better position the proposed demonstrator within the current state of the art, Table 4 compares representative studies on automated wood-defect detection. Unlike recent research, which has primarily focused on improving classification performance through increasingly sophisticated deep-learning architectures trained on large image datasets, the present work addresses the laboratory validation of a complete embedded computer-vision system intended for early-stage technology maturation. Therefore, the comparison emphasises methodological characteristics and research focus rather than direct numerical performance, since the reported results were obtained under substantially different experimental conditions.

TABLE 4
POSITIONING OF THE PROPOSED EMBEDDED COMPUTER-VISION DEMONSTRATOR WITH RESPECT TO REPRESENTATIVE STUDIES ON AUTOMATED WOOD-DEFECT DETECTION

Reference	Method	Experimental data	Embedded implementation	Research focus
Hu et al. [5]	GLCM-based texture analysis	Knot image dataset	No	Classical machine-vision defect detection
Urbonas et al. [18]	Faster R-CNN	Veneer image dataset	No	Deep-learning veneer defect detection
Gao et al. [19]	ResNet-34 (transfer learning)	Knot image dataset	No	CNN-based knot classification
An et al. [4]	CWB-YOLOv8	Multi-species wood image dataset	No	Real-time wood defect detection
Lin et al. [17]	WDNET-YOLO	Structural timber dataset	No	Structural timber inspection
Wang et al. [13]	Improved YOLOv8	Wood surface dataset	No	Surface defect detection
This work	**Random Forest + embedded vision system**	**24 industrial beech boards (192 image regions)**	**Yes**	**TRL4 validation of an integrated embedded inspection system**

As shown in Table 4, most published studies concentrate on algorithmic optimisation using large annotated datasets and high-performance computing platforms, whereas comparatively little attention has been paid to the experimental validation of complete embedded inspection systems. The proposed demonstrator addresses this gap by integrating image acquisition, pre-processing, feature extraction and classification into a low-cost embedded architecture validated at Technology Readiness Level 4, thereby establishing the technological foundation for subsequent industrial development within the SMARTWOOD-AI platform.

An important aspect of the proposed methodology is the adoption of board-level data partitioning, which prevents information leakage between the training and validation sets. In wood inspection, several image patches frequently originate from the same physical board. Splitting datasets at image level may therefore allow visually similar regions from the same specimen to appear simultaneously in both subsets, artificially increasing the reported performance. By performing the split at the physical-board level, the present study provides a more conservative and realistic estimate of classifier generalisation, which is particularly important for small experimental datasets intended for technology validation rather than algorithm benchmarking.

Another noteworthy outcome concerns the feature-importance analysis. The dominant contribution of GLCM- and LBP-based descriptors confirms that surface texture remains the most discriminative source of information for identifying natural wood defects. This observation is fully consistent with previous studies demonstrating that texture descriptors retain considerable discriminative power even in the presence of heterogeneous wood anatomy [5-9]. At the same time, the comparatively smaller contribution of HSV colour features suggests that chromatic information alone is insufficient for distinguishing defects from natural variations in beech wood, reinforcing the importance of combining complementary descriptor families.

3.5.1 Classification Error Analysis

The detailed analysis of classification errors provides further insight into the behaviour of the proposed model. The classifier achieved high performance for the defect class (precision = 0.84, recall = 0.93), indicating that the system successfully identifies the majority of surfaces containing visible defects. Conversely, the performance for the nodefect class (precision = 0.78, recall = 0.58) is substantially lower. This imbalance is largely explained by the experimental dataset itself: although segmentation increased the number of image samples, only two physical boards without defects were available for training, limiting the actual variability represented by this class.

A detailed investigation of the misclassified board01 illustrates this phenomenon. The board was visually confirmed to be free of natural defects but was predominantly classified as defective. An initial hypothesis attributed this behaviour to chromatic characteristics, assuming that its lighter colour might resemble discolouration defects. However, comparison of the HSV feature distributions showed that the colour statistics of board01 were considerably closer to those of the nodefect class than to those of defective boards. Consequently, chromatic information alone cannot explain the observed misclassification.

Instead, detailed visual inspection revealed superficial circular-saw machining traces accompanied by local fibre disruption. Although these artefacts do not represent natural wood defects, they modify the local texture sufficiently to generate responses similar to those produced by knots or cracks in GLCM- and LBP-based descriptors. This interpretation is consistent with previous reports highlighting the influence of wood texture and machining artefacts on automated inspection performance [17]. The results therefore suggest that the current classifier is more sensitive to structural texture variations than to colour differences, an observation that agrees with the feature-importance analysis presented in Section 3.4.

Several strategies could be considered to mitigate the influence of machining artefacts in future work. These include (i) additional preprocessing steps such as directional filtering to reduce the prominence of saw marks, (ii) the use of orientation-invariant texture features that are less sensitive to directional patterns, and (iii) the incorporation of larger spatial context windows to distinguish between local machining traces and genuine defect patterns. Such approaches would likely improve the specificity of the classifier without compromising its sensitivity to actual defects.

From an industrial perspective, the observed error distribution is not necessarily undesirable. The classifier favours high sensitivity to defects, producing relatively few false negatives while accepting a higher number of false positives. Since a sound board incorrectly classified as defective can subsequently be verified by an operator, whereas an undetected defect may compromise downstream processing and product quality, this operating point represents a reasonable compromise for an early-stage inspection system.

3.5.2 Limitations

The present results should be interpreted within the context of a laboratory demonstrator developed for TRL4 validation. Several limitations must therefore be acknowledged.

First, the experimental dataset remains relatively small and exhibits a pronounced class imbalance, with only two physically defect-free boards. Although image segmentation increased the number of training samples, it could not compensate for the limited diversity of the underlying material. This limitation is particularly relevant for the nodefect class, where the small sample size may not adequately represent the full range of natural variation in defect-free beech wood.

Second, the feature-based Random Forest approach deliberately prioritises interpretability and computational efficiency over maximum predictive performance. While hand-crafted descriptors proved suitable for validating the technological concept, they cannot fully capture the complex hierarchical image representations learned automatically by modern deep-learning architectures. Consequently, the reliable identification of subtle defects such as micro-cracks, incipient knots or irregular grain distortions will likely require CNN- or YOLO-based models trained on substantially larger datasets [16].

Third, several uncontrolled industrial artefacts—including machining marks, chalk annotations, glove traces and transport-induced surface modifications—were observed during the experimental campaign. Although their influence was documented qualitatively, their quantitative impact on classifier performance was not analysed separately. Future investigations should therefore include systematic annotation of such artefacts in order to distinguish manufacturing traces from genuine wood defects.

Finally, the present work validates only the computer-vision module of the SMARTWOOD-AI architecture. The digital-twin framework, multi-objective cutting optimisation and closed-loop decision system constitute subsequent development stages and were intentionally excluded from the scope of the present TRL4 validation.

3.6 Industrial Implications and Integration into SMARTWOOD-AI

The proposed demonstrator should be regarded as the first operational building block of the SMARTWOOD-AI platform rather than as a complete industrial inspection solution. Its primary contribution is the experimental validation of a fully integrated embedded vision system capable of generating objective defect information under controlled laboratory conditions.

Within the planned SMARTWOOD-AI architecture, these defect maps will provide the input for digital-twin models representing individual timber pieces, enabling multi-objective optimisation of cutting strategies and closed-loop decision support [30-33]. This development trajectory is fully aligned with recent research advocating the integration of computer vision, digital twins and intelligent production planning for next-generation sawmills [28, 31-33].

The technological choices adopted in the present work also support this progression. The use of a low-cost embedded platform facilitates future industrial deployment, while the interpretable Random Forest model establishes a transparent experimental baseline against which more sophisticated CNN- and YOLO-based approaches can be objectively evaluated. As the experimental database expands with images collected under real industrial conditions, the current feature-based methodology will gradually evolve towards multiclass detection and localisation models capable of identifying individual defect types, ultimately supporting the progression from laboratory validation (TRL4) to industrial demonstration (TRL7) and technology transfer (TRL9).

IV. CONCLUSION

This paper presented the laboratory validation of a low-cost embedded computer-vision demonstrator for the automated detection of defects in beech sawn timber, developed as the first functional module of the SMARTWOOD-AI platform. The study demonstrated that an integrated inspection chain combining embedded image acquisition, image preprocessing, hand-crafted feature extraction and Random Forest classification can successfully operate under controlled laboratory conditions, thereby fulfilling the requirements associated with Technology Readiness Level 4 (TRL4 – Technology validated in laboratory).

The experimental evaluation confirmed the technical feasibility of the proposed approach. The demonstrator achieved 82.5% accuracy on an independent validation set, while maintaining high sensitivity for defect detection (precision = 0.84, recall = 0.93 for the defect class). Although the experimental dataset was intentionally limited and highly imbalanced, the adopted methodology—including board-level data partitioning, feature-importance analysis and critical error evaluation—provided a realistic assessment of the system performance and established a robust experimental baseline for future developments.

Beyond the reported classification performance, the principal contribution of this work lies in the validation of an embedded inspection architecture rather than in the optimisation of a single machine-learning algorithm. The proposed system demonstrates that interpretable classical computer-vision techniques remain a practical solution during the early stages of technology maturation, particularly when limited datasets, low computational requirements and transparent decision-making are important design constraints. Furthermore, the integration of a Raspberry Pi 4 with a Sony IMX500 intelligent vision sensor represents a promising embedded architecture for future industrial deployment of AI-assisted timber inspection systems.

The findings also identify the main challenges that must be addressed before industrial implementation. These include increasing the diversity and size of the experimental dataset, improving robustness against machining artefacts and natural texture variability, and extending the methodology from binary classification towards multiclass defect detection and localisation using deep-learning models.

Finally, this work establishes the technological foundation for the subsequent development of the SMARTWOOD-AI platform. The validated computer-vision module will provide objective defect information to future digital-twin models, multi-objective cutting optimisation algorithms and closed-loop decision-support mechanisms. The present work therefore provides not only a validated embedded computer-vision demonstrator but also a reproducible experimental framework that can support the progressive maturation of AI-assisted timber inspection systems from laboratory validation (TRL4) towards industrial deployment (TRL9).

CONFLICT OF INTEREST

The authors declare no conflicts of interest.

REFERENCES

[1] An, H., Liang, Z., Qin, M., Huang, Y., Xiong, F., & Zeng, G. (2024). Wood defect detection based on the CWB-YOLOv8 algorithm. *Journal of Wood Science*, *70*(1), Article 26. https://doi.org/10.1186/s10086-024-02139-z

[2] Capogrosso, L., Bonazzi, P., Hoxhaj, L., & Magno, M. (2026). *Exploiting in-sensor computing for energy-efficient earth observation*.

[3] Casas, G. G., Ismail, Z. H., & Leite, H. G. (2026). Computer vision, machine learning, and deep learning for wood and timber products: A Scopus-based bibliometric and systematic mapping review (1983–2026, early access). *Forests*, *17*(1), Article 112. https://www.mdpi.com/1999-4907/17/1/112

[4] Chabanet, S., Bril El-Haouzi, H., Morin, M., Gaudreault, J., & Thomas, P. (2023). Toward digital twins for sawmill production planning and control: Benefits, opportunities, and challenges. *International Journal of Production Research*, *61*(7), 2190-2213. https://doi.org/10.1080/00207543.2022.2068086

[5] Chabanet, S., El Haouzi, H., & Thomas, P. (2023). Toward a sawmill digital shadow based on coupled simulation and supervised learning models. In *Proceedings* (pp. 59-70).

[6] Deniz, M., Bogrekci, I., & Demircioglu, P. (2025). Real-time detection of hole-type defects on industrial components using Raspberry Pi 5. *Applied System Innovation*, *8*(4), Article 89. https://www.mdpi.com/2571-5577/8/4/89

[7] Fan, C., Zhuang, Z., Liu, Y., Yang, Y., Zhou, H., & Wang, X. (2024). Bilateral defect cutting strategy for sawn timber based on artificial intelligence defect detection model. *Sensors*, *24*(20), Article 6697. https://www.mdpi.com/1424-8220/24/20/6697

[8] Fredriksson, M. (2015). Optimizing sawing of boards for furniture production using CT log scanning. *Journal of Wood Science*, *61*(5), 474-480. https://doi.org/10.1007/s10086-015-1500-0

[9] Gao, M., Qi, D., Mu, H., & Chen, J. (2021). A transfer residual neural network based on ResNet-34 for detection of wood knot defects. *Forests*, *12*(2), Article 212. https://www.mdpi.com/1999-4907/12/2/212

[10] Han, S., Jiang, X., & Wu, Z. (2023). An improved YOLOv5 algorithm for wood defect detection based on attention. *IEEE Access*, *11*, 71800-71810. https://doi.org/10.1109/ACCESS.2023.3293864

[11] He, T., Liu, Y., Xu, C., Zhou, X., Hu, Z., & Fan, J. (2019). A fully convolutional neural network for wood defect location and identification. *IEEE Access*, *7*, 123453-123462. https://doi.org/10.1109/ACCESS.2019.2937461

[12] He, T., Liu, Y., Yu, Y., Zhao, Q., & Hu, Z. (2020). Application of deep convolutional neural network on feature extraction and detection of wood defects. *Measurement*, *152*, Article 107357. https://doi.org/10.1016/j.measurement.2019.107357

[13] Hittawe, M. M., Sidibé, D., & Mériaudeau, F. (2015). A machine vision based approach for timber knots detection. In *The International Conference on Quality Control by Artificial Vision 2015*. SPIE.

[14] Hu, C., Min, X., Yun, H., Wang, T., & Zhang, S. (2011). Automatic detection of sound knots and loose knots on sugi using gray level co-occurrence matrix parameters. *Annals of Forest Science*, *68*(6), 1077. https://doi.org/10.1007/s13595-011-0123-x

[15] Hu, J., Song, W., Zhang, W., Zhao, Y., & Yilmaz, A. (2019). Deep learning for use in lumber classification tasks. *Wood Science and Technology*, *53*(2), 505-517. https://doi.org/10.1007/s00226-019-01086-z

[16] Hwang, S.-W., Lee, T., Kim, H., Chung, H., Choi, J. G., & Yeo, H. (2021). Classification of wood knots using artificial neural networks with texture and local feature-based image descriptors. *Holzforschung*, *76*(1), 1-13. https://doi.org/10.1515/hf-2021-0051

[17] Ji, M., Zhang, W., Diao, X., Wang, G., & Miao, H. (2023). Intelligent automation manufacturing for Betula solid timber based on machine vision detection and optimization grading system applied to building materials. *Forests*, *14*(7), Article 1510. https://www.mdpi.com/1999-4907/14/7/1510

[18] Ji, M., Zhang, W., Han, J.-k., Miao, H., Diao, X.-l., & Wang, G.-f. (2024). A deep learning-based algorithm for online detection of small target defects in large-size sawn timber. *Industrial Crops and Products*, *222*, Article 119671. https://doi.org/10.1016/j.indcrop.2024.119671

[19] Kamal, K., Qayyum, R., Mathavan, S., & Zafar, T. (2017). Wood defects classification using laws texture energy measures and supervised learning approach. *Advanced Engineering Informatics*, *34*, 125-135. https://doi.org/10.1016/j.aei.2017.09.007

[20] Kılıç, K., Kılıç, K., Doğru, İ., & Özcan, U. (2025). WD Detector: Deep learning-based hybrid sensor design for wood defect detection. *European Journal of Wood and Wood Products*, *83*. https://doi.org/10.1007/s00107-025-02211-5

[21] Li, R., Zhong, S., & Yang, X. (2025). Wood panel defect detection based on improved YOLOv8n. *BioResources*, *20*, 2556-2573. https://doi.org/10.15376/biores.20.2.2556-2573

[22] Lin, X., et al. (2025). WDNET-YOLO: Enhanced deep learning for structural timber defect detection to improve building safety and reliability. *Buildings*, *15*(13), Article 2281. https://www.mdpi.com/2075-5309/15/13/2281

[23] Lukovic, M., et al. (2024). Probing the complexity of wood with computer vision: From pixels to properties. *Journal of The Royal Society Interface*, *21*(213). https://doi.org/10.1098/rsif.2023.0492

[24] Nasir, V., Rahimi, S., Mohammadpanah, A., Hansen, E., & Sassani, F. (2024). Intelligent lumber production (Sawmill 4.0): Opportunities, challenges, and pathways to adoption. In M.-R. Alam & M. Fathi (Eds.), *Integrated systems: Data driven engineering* (pp. 213-231). Springer Nature Switzerland.

[25] Okano, M. T., Lopes, W. A. C., Ruggero, S. M., Vendrametto, O., & Fernandes, J. C. L. (2025). Edge AI for industrial visual inspection: YOLOv8-based visual conformity detection using Raspberry Pi. *Algorithms*, *18*(8), Article 510. https://www.mdpi.com/1999-4893/18/8/510

[26] Ozkan, M., Ozcan, C., & Gökmen, S. (2025). Deep learning based defect detection and quality classification on lamella pieces used in solid wood panel production. *Operations Research Forum*, *6*. https://doi.org/10.1007/s43069-025-00534-w

[27] Popa, S. E., Puiu, P. G., Andrioaia, D. A., Grigore, R. M., & Avădanei, R. L. (2026). Indirect estimation of absorbed infrared LED radiant power using contactless thermal sensing. *Sensors*, *26*(13), Article 4055. https://www.mdpi.com/1424-8220/26/13/4055

[28] Urbonas, A., Raudonis, V., Maskeliūnas, R., & Damaševičius, R. (2019). Automated identification of wood veneer surface defects using faster region-based convolutional neural network with data augmentation and transfer learning. *Applied Sciences*, *9*(22), Article 4898. https://www.mdpi.com/2076-3417/9/22/4898

[29] Wang, B., Wang, R., Chen, Y., Yang, C., Teng, X., & Sun, P. (2025). FDD-YOLO: A novel detection model for detecting surface defects in wood. *Forests*, *16*(2), Article 308. https://www.mdpi.com/1999-4907/16/2/308

[30] Wang, R., Chen, Y., Liang, F., Wang, B., Mou, X., & Zhang, G. (2024). BPN-YOLO: A novel method for wood defect detection based on YOLOv7. *Forests*, *15*(7), Article 1096. https://www.mdpi.com/1999-4907/15/7/1096

[31] Wang, R., Liang, F., Wang, B., Zhang, G., Chen, Y., & Mou, X. (2024). An efficient and accurate surface defect detection method for wood based on improved YOLOv8. *Forests*, *15*(7), Article 1176. https://www.mdpi.com/1999-4907/15/7/1176

[32] Wu, L., et al. (2026). A lumber surface defect detection network integrating deformable convolution and multi-scale attention. *Forests*, *17*(7), Article 839. https://www.mdpi.com/1999-4907/17/7/839

[33] Xi, H., Wang, R., Liang, F., Chen, Y., Zhang, G., & Wang, B. (2024). SiM-YOLO: A wood surface defect detection method based on the improved YOLOv8. *Coatings*, *14*(8), Article 1001. https://www.mdpi.com/2079-6412/14/8/1001.