An Attention-Enhanced Deep Learning Framework for Automated Quality Grading of Ribbed Smoked Rubber Sheets

D. P. Sumalatha
Department of Computer Science and Engineering, KVG College of Engineering, Sullia, Karnataka, India
suma_devajana@yahoo.com (corresponding author)

M. L. Smitha
Department of Computer Science and Engineering, KVG College of Engineering, Sullia, Karnataka, India
smithaml.urubail@gmail.com

U. J. Ujwal
Department of Computer Science and Engineering, KVG College of Engineering, Sullia, Karnataka, India
ujwalu@yahoo.com

Received: 6 November 2025 | Revised: 12 December 2025 | Accepted: 29 December 2025

ABSTRACT

The grading of Ribbed Smoked Rubber Sheets (RSS) plays a critical role in ensuring consistency, quality, and market value within the natural rubber industry. Traditional manual inspection methods are subjective, time-consuming, and prone to human error. With the increasing demand for standardized systems, there is a need for automating the RSS grading process using vision-based techniques. This study proposes an Attention-Enhanced Deep Learning Framework (AEDL-RSS) that combines convolutional feature extraction with spatial and channel attention modules to enhance robustness against lighting, color variation, and surface defects. The model integrates an attention-augmented Convolutional Neural Network (CNN) backbone and a lightweight classifier for real-time deployment. The experimental results demonstrate that the proposed system outperforms the baseline VGG16 model, achieving a classification accuracy of 95% on the test set. The system provides a reliable, scalable, and objective alternative to manual inspection, offering potential for real-time industrial deployment and smart factory integration in the rubber processing sector.

Keywords-attention mechanism; CBAM; CNN; deep learning; K-Fold cross-validation; rubber sheet grading

I. INTRODUCTION

RSS are among the most widely traded forms of processed natural rubber. Their grading determines market pricing, product suitability, and export standards. Manual grading is currently performed by trained experts who visually inspect each sheet based on parameters such as color uniformity, surface cleanliness, texture pattern, and presence of bubbles or blemishes. Accurate grading affects the market value and global trading standards. Therefore, precise grading is very crucial for maintaining quality standards and ensuring fair pricing.

Over the past decade, deep learning has gained acceptance in agricultural product assessment due to its ability of learning complex features. Advanced models, such as CNNs, have been applied successfully in agricultural grading, fruit quality assessment, and industrial surface inspection through their ability to extract hierarchical features [1]. The introduction of attention mechanisms has further improved the fine-grained classification by using spatial and channel attention mechanisms. Authors in [2, 3] showed the attention enhancement feature interpretability and model focus, especially in complex visual datasets. Such mechanisms are useful for RSS grading, where subtle features are crucial.

Earlier attempts of RSS grading utilized traditional techniques, which include color-based thresholding [4], handcrafted feature extraction [4], statistical texture measures, and neural networks such as Learning Vector Quantization

(LVQ) [5] or perceptron network [6]. An earlier system, the Ribbed Smoked Sheet Grading System (RSSGS) [7], also employed image processing techniques; however, lacked robustness under illumination. Although these approaches exhibited feasibility, they suffer from poor generalization when exposed to illumination, camera quality, and orientation in rubber sheets. Agricultural grading systems, such as pineapple or fruit grading [8, 9], and CNN-based agricultural grading systems [10, 11], show that deep learning provides improved performance over other traditional approaches. However, automated grading of RSS remains underexplored, with no earlier works using attention modules, such as Convolutional Block Attention Module (CBAM) [12], or cross-validation techniques [13].

To address these challenges, this work proposes an AEDL-RSS that integrates an attention-augmented CNN with domain-aware preprocessing. The proposed model integrates VGG16 backbone with CBAM [12, 14] for fine-grained feature extraction and focuses on relevant features, thereby improving recognition of fine-grained defects. The model is evaluated using 5-fold cross-validation, using an unbiased accuracy estimation proposed in [13]. This study is the first to integrate CBAM into rubber sheet grading for fine-grained feature extraction, employed with K-Fold cross-validation for statistically robust evaluation, and introduces a domain-specific preprocessing and patch-based dataset strategy addressing the texture variability of real RSS sheets. It proposes a complete attention-enhanced RSS grading framework (AEDL-RSS) with Focal loss to handle imbalance and experimental comparison covering modern architectures and ablation studies. These contributions represent practical advancements beyond prior RSS research.

The main contributions of this study are:

- A domain-specific preprocessing pipeline designed to mitigate lighting and color inconsistencies in RSS images.
- An attention-enhanced CNN architecture combining CBAM is integrated with the base VGG16 model for context modeling.
- A comprehensive experimental evaluation, including ablation studies, comparison with modern architectures (ResNet50, EfficientNet-B0, MobileNetV2), computational efficiency analysis, and confusion matrix–based error assessment.
- A robust and deployable RSS grading framework, demonstrating 95.05% accuracy, offering a practical solution for manual grading.

II. METHODOLOGY

A. Dataset

A custom dataset of high-resolution RSS real-world images was collected in 2025 from a certified rubber society in Kerala, India. All images were captured outdoors with bright lighting conditions, using a Canon DSLR camera. A consistent capture distance and frontal orientation were mentioned during image acquisition. The dataset covers all five grades (1-5) in accordance with industry standards.

A total of 88 full-sheet images were collected. Although the number of samples is limited, this is common in agricultural quality-assessment studies where:

- Data collection requires expert-verified grading.
- Access to industry-graded sheets is restricted.
- The physical size of RSS sheets limits the quantity that can be captured in a single session.

All images were graded by certified rubber expert graders of the Rubber Society during dataset collection. Additional inter-rater reliability was not required as the annotations were provided by domain experts following standard RSS grading protocols. To increase the data volume, all images were converted into patches.

B. Patch Extraction

For dataset scalability and capturing localized texture variations, a non-overlapping patch of size 224 ×224 was extracted from each image. The training set comprised 4,088 patches, while the testing set included 1,025 patches. Patches retain fine-grained texture, such as tiny bubbles and cracks. The detailed data split is presented in Table I.

TABLE I. PATCH DISTRIBUTION ACROSS TRAINING AND TEST SPLITS

Grade	Train	Test	Total
Grade 1	1036	259	1295
Grade 2	1326	332	1658
Grade 3	944	237	1181
Grade 4	572	144	716
Grade 5	210	53	263

C. Offline Data Augmentation

To address class imbalance, offline synthetic augmentation was applied, and all five grades were balanced to 1888 images.

The augmentation techniques used include rotation (±20°), horizontal flip, zoom (0.2), width/height shift (0.1), and brightness adjustment (0.8–1.2).

D. Sample Images Per Grade

Figure 1 illustrates representative original RSS sheet images of all five grades. These samples show the natural visual differences across grades, including surface texture, color uniformity, bubble density, and any other defects that appear severe. Such visual differences confirm the need for patch extraction and a deep learning approach to capture fine-grained texture patterns.

E. Model Architecture

The proposed architecture is built on the VGG16 CNN, a deep learning model with proven records in image classification tasks due to its structured layer design. CBAM is integrated into the backbone of the model to enhance its ability to focus on spatial and channel-wise features. The base VGG16 model, which is pre-trained on ImageNet, is adopted by allowing fully connected layers customized by implementing the following modifications:

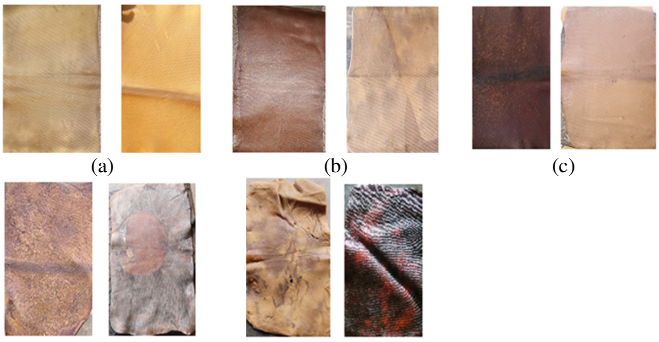

Fig. 1. Sample original RSS images: (a) Grade 1, (b) Grade 2, (c) Grade 3, (d) Grade 4, (e) Grade 5.

1) CBAM Integration

CBAM modules are implemented after the fifth block of the convolutional layer in VGG16. CBAM provides spatial and channel attention, which enables the model to focus only on relevant features rather than irrelevant regions. Channel Attention and Spatial Attention are calculated as shown in:

$$M_c\,(F) = \sigma(\mathrm{MLP}(\mathrm{AvgPool}(F)) + \mathrm{MLP}(\mathrm{MaxPool}(F)) \quad (1)$$

$$M_s\,(F) = \sigma(f_{7\times7}([AvgPool(F); MaxPool(F)])) \qquad (2)$$

2) Global Average Pooling (GAP)

GAP helps reduce the overfitting problems by maintaining spatial features.

3) Fully Connected Layers

Dense layers with an activation function called Rectified Linear Unit (ReLu) were utilized to introduce non-linearity and enable the model in learning complex features.

4) Softmax

The final layer of the dense layer uses Softmax activation to classify the output into five different grades utilizing probability density. Figure 2 illustrates the model architecture, which is trained using Focal loss as a loss function, Adam optimizer, learning rate scheduling, and early stopping depending on validation accuracy.

F. K-Fold Cross Validation

To prevent overfitting and improve generalization, K-Fold Cross Validation was adopted, which is a robust evaluation technique. The dataset was split into K-sized equal folds with *K* = 5 and was run for 10 epochs per fold. On/During every iteration, one of the folds is used for validation, and the remaining four folds are the training set. Performance metrics, such as accuracy, precision, recall, and F1-score, are averaged from all the folds to obtain an impartial estimation of the model's performance. This strategy helps in increasing the use of limited labeled data, detecting overfitting, and variance across splits.

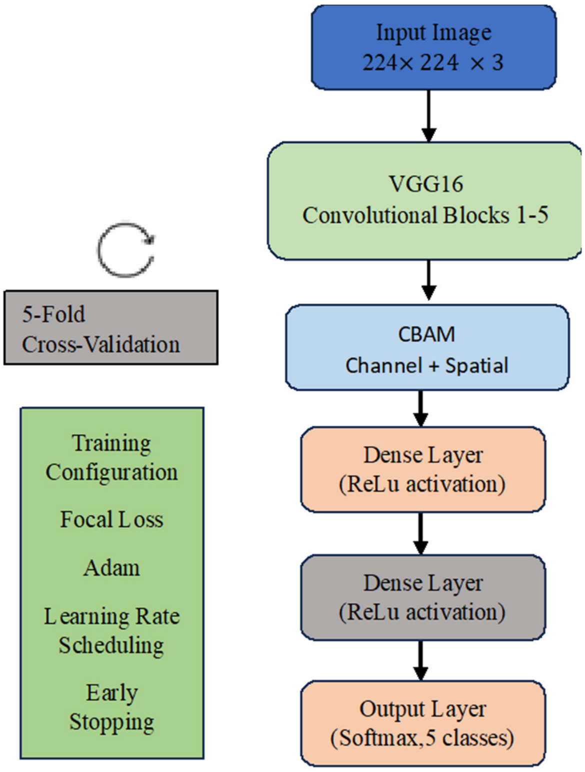

Fig. 2. Model architecture.

a) Workflow

The dataset was divided into 5 equal folds. For every iteration 4 folds were used for training, while the remaining fold was used for validation. This process was repeated five times, changing validation for each fold. The complete workflow diagram is portrayed in Figure 3.

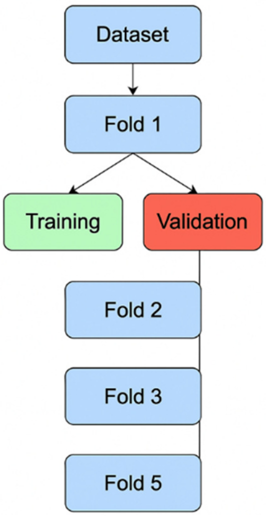

Fig. 3. K-Fold cross-validation workflow.

III. EXPERIMENTS AND RESULTS

A. Setup

The experiments were conducted using Google Colab Pro with NVIDIA Tesla T4 GPUs. Adam Optimizer was employed for optimization with a learning rate of 1×10^{-4}. Class imbalance across all five grades of rubber sheets was handled using Focal loss. The model was trained for 10 epochs with a batch size of 32. Callbacks, such as EarlyStopping (patience = 5) and ModelCheckpoint, were used to save the best model during training.

1) Handling Class Imbalance:

Class imbalance across all five grades was handled using Focal loss with $\gamma = 2.0$. Focal loss is calculated as:

$$FL = -w_i(1 - {p_t}^{\gamma}) \log (p_t) \quad (3)$$

Class Weights were computed for each fold using: `class_weight.compute_class_weight`.

2) Hyperparameter Selection

Learning rates of 1×10^{-3} to 1×10^{-5} and batch sizes of 16 and 32 were explored. The chosen configuration provided the best balance between accuracy and speed.

3) Reproducibility

Random seeds were fixed for NumPy and TensorFlow (seed = 42) to ensure reproducibility. All experiments were executed using TensorFlow 2.x.

4) Training Duration

The model was trained for 10 epochs per fold for the 5-fold cross-validation to ensure consistent evaluation across folds and maintain computational feasibility.

For the baseline and other comparative models, training was performed for 30 epochs, allowing them to connect in one single-split training.

B. Performance Metrics

The trained model was evaluated utilizing standard classification metrics, namely accuracy, precision, recall, and F1-score, which were calculated for each rubber sheet grade and averaged. Metrics of training, validation, and test datasets were recorded for the comparison of generalization.

1) Classification Report

The detailed Classification performance of VGG16+CBAM with K-fold cross-validation is summarized in Table II.

TABLE II. SUMMARY OF PERFORMANCE METRICS

Class	**Precision (%)**	**Recall (%)**	**F1-score (%)**
Grade 1	99.0 ± 1.0	98.0 ± 0.0	99.0 ± 0.0
Grade 2	99.0 ± 1.0	96.0 ± 2.0	97.0 ± 0.0
Grade 3	96.0 ± 1.0	93.0 ± 3.0	94.0 ± 2.0
Grade 4	88.0 ± 5.0	89.0 ± 4.0	89.0 ± 2.0
Grade 5	78.0 ± 6.0	95.0 ± 3.0	85.0 ± 3.0
Overall accuracy	95.05 ± 0.90		

The proposed model achieved an overall accuracy of 95.05%, with a macro-average F1-score of 0.93 and a weighted average F1-score of 0.95 across all grades. The model performed well in Grade 1 and Grade 2 with an F1-score >0.97, while its performance on Grade 5 was relatively lower, obtaining an F1-score of 0.85. This can be attributed to the limited number of samples in Grade 5, leading to class imbalance. However, the weighted F1-score of 0.95 indicates robust overall performance.

2) Confusion Matrix Visualization

Figure 4 displays the confusion matrix for the 5-fold cross-validation with the distribution of true versus predicted labels. The confusion matrix shows a detailed view of the model's performance. Samples that appeared diagonally indicate a high proportion of correct classifications across all RSS grades. Misclassifications occur rarely between visually similar neighboring grades, indicating the natural overlap in texture and color patterns observed in real RSS sheets. The overall matrix demonstrates that the model captures discriminative features effectively through the attention mechanisms, resulting in reliable grading of RSS sheets.

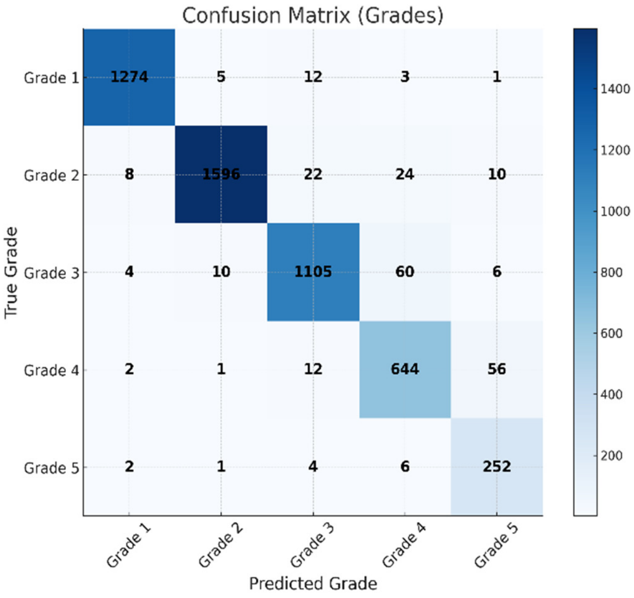

Fig. 4. Confusion matrix for the 5-fold cross-validation showing true and predicted labels.

3) CBAM Visualization

Figure 5 shows the CBAM spatial attention visualizations for Grade 2 and Grade 5, with the original image, CBAM spatial attention map, and overlay. Warmer colors indicate regions with higher attention weights. CBAM visualizations were generated using the final CBAM block (CBAM5), as it provides the most semantically meaningful spatial attention maps for model interpretation.

4) Error Analysis Report

The classification report demonstrates that the greatest misclassification is observed between Grade 3 and Grade 4, and between Grade 4 and Grade 5, where visual characteristics, such as texture density and discoloration, occur significantly.

These grades show overlapping surface textures, making the boundary between moderate and severe defects visually obscure. Grade 5 exhibits the lowest F1-score compared to other grades due to limited training samples and high intra-class variance. CBAM further demonstrates diffuse focus regions for Grade 5, supporting the observed error patterns.

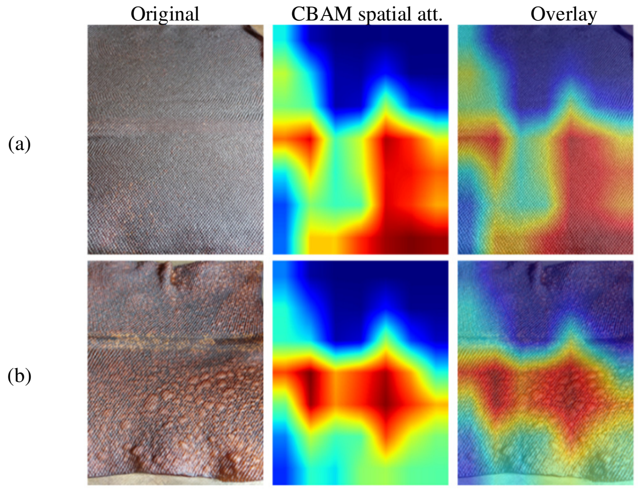

Fig. 5. CBAM spatial attention visualizations for: (a) Grade 2, and (b) Grade 5.

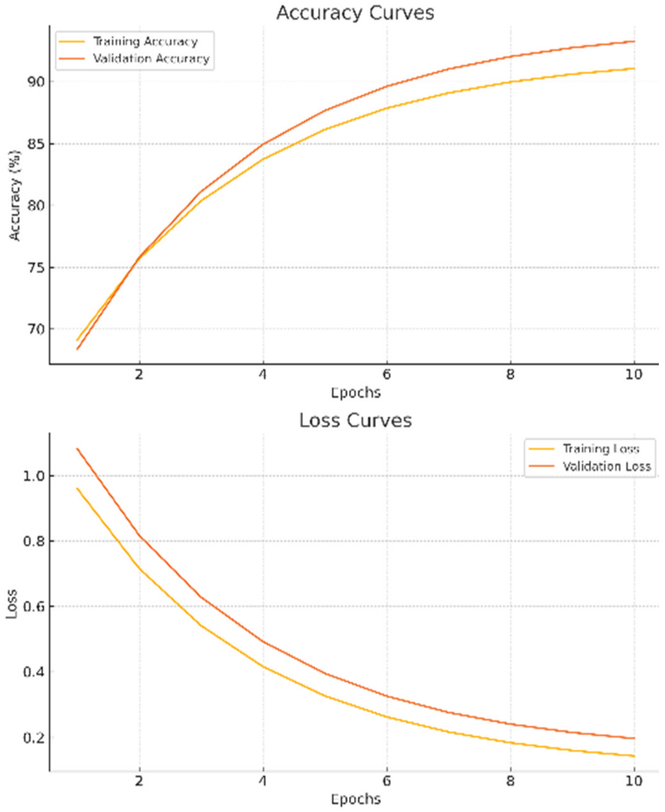

Fig. 6. Training and validation accuracy/loss curves for the proposed model.

5) *Learning Curves*

The curve in Figure 6 shows stable convergence with minimal overfitting, indicating good generalization across all folds.

C. *Ablation Study*

To evaluate the effectiveness of the proposed model, the latter was compared with baseline architectures, such as VGG16, VGG16+K-Fold, VGG16+CBAM, and VGG16+CBAM+K-Fold, trained under identical conditions.

TABLE III. VALIDATION RESULTS (MODEL SELECTION STAGE)

Model	Validation accuracy (%)
VGG16(baseline)	91.99
VGG16 + K-Fold	92.40
VGG16 + CBAM	96.55
VGG16 + CBAM +K-Fold (proposed)	94.57

All accuracy values reported in Table III correspond to the mean validation accuracy across 5 folds. VGG16+CBAM+K-Fold achieved an accuracy of 94.57±1.02 across all 5 folds, demonstrating stable performance across all splits.

The ablation study demonstrates that the addition of K-Fold validation has only a marginal improvement, indicating that the VGG16 model benefits slightly from better training-validation separation, but feature learning is not improved intrinsically. Moreover, the introduction of CBAM boosts accuracy remarkably (+4.56%), exhibiting that channel and spatial attention mechanisms help in learning discriminative features critical for rubber sheet grading. The combination of CBAM with K-Fold shows a slight reduction in accuracy compared to VGG16+CBAM, as each fold was trained for only 10 epochs, but shows much lower variance, which indicates strong generalization and robustness.

Although VGG16+CBAM achieves higher validation accuracy in a single split, the proposed VGG16+CBAM+K-Fold model balances between accuracy and reliability, making it the best for real-world implementation.

D. *Computational Analysis*

CBAM introduces negligible computational overhead (+0.05 GFLOPs) while improving fine-grained feature learning. Inference time remains within 3-7 ms per image.

TABLE IV. MODEL SIZE, FLOPS, AND INFERENCE TIME COMPARISON

Model	FLOPs	Inference time (ms/image)
VGG16(baseline)	15.50	3-6
VGG16 + GAP + Dense	15.51	3-6
VGG16 + CBAM	15.55	4-7

The VGG16-CBAM model requires approximately 64 MB for storing trainable parameters and around 450 MB of GPU memory during inference, rendering it deployable on standard 8-12 GB GPU systems. Since K-Fold is a training strategy and does not alter the inference architecture, FLOPs and Inference time are reported only for the VGG16-CBAM model.

E. Test Results

The performance comparison, as presented in Table V, demonstrates that CBAM integration with VGG16 improves the model's capability in discriminating fine-grained features, with improved test accuracy from 91.71% to 95.05%. The proposed VGG16+CBAM+K-fold model achieves stable performance with a mean test accuracy of 95.05%±0.90, demonstrating its reliability and robustness. Although the average accuracy is slightly lower than the VGG16+CBAM model, the proposed model achieves better generalization with consistent precision, recall, and F1-scores. This proves that the model's reliability is enhanced by consolidating both cross-validation and attention mechanisms in real-world rubber sheet grading.

F. Discussion

The CBAM module was used for better feature extraction, focusing on key regions of the image, improving fine-grained discrimination. The use of K-fold cross-validation showed coherent accuracy in validation across all folds, proving that the proposed method is robust. Even though slight-overfitting was observed in the CBAM implementation alone, K-Fold integration aided in alleviating this issue with increased performance. The confusion matrix analysis and CBAM visualizations show strong classification for most grades, mainly Grade 1 and Grade 2, with F1-scores above 0.97. Nevertheless, Grade 5 exhibited lower performance compared to the F1-score of 0.85, likely due to the smaller number of samples indicating class imbalance. These findings suggest potential usage in automation grading.

1) Comparison with Modern Architecture

Table VI compares the proposed model with MobileNet, EfficientNet, and ResNet architectures. While all models performed well, the proposed VGG16-CBAM model achieved the highest overall accuracy. MobileNet, being a lightweight model, did not perform well with global features, whereas EfficientNet-B0 and ResNet-50 showed low performance. The higher accuracy of the VGG16-CBAM model highlights the benefits of integrating attention mechanisms for subtle defect identification in RSS sheets. The results of other architectures are taken directly from published literature and are not reproduced using the proposed dataset.

TABLE V. TEST RESULTS (FINAL EVALUATION STAGE)

Model	Test accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
VGG16	91.71	88	93	90
VGG16+ CBAM	99.02	98	98	98
VGG16+ CBAM+ K-Fold (proposed)	95.05±0.90	92	94	92

TABLE VI. COMPARISON WITH MODERN ARCHITECTURES

Architecture	Parameters (M)	Test accuracy (%)	Notes
MobileNet	3.4M	92.10	Lightweight, weaker texture capture
EfficientNet-B0	5.3M	94.00	Good but unstable on fine textures
ResNet-50	25.6M	93.40	Heavy model, higher inference time
VGG16+CBAM+K-Fold (proposed)	16.8M	95.04	Best fine texture recognition

Overall, the experimental results confirm the performance of the VGG16-CBAM framework for automatic grading of ribbed smoke rubber sheets with high accuracy, strong generalization, and reliability for real-world implementation, reducing the subjectivity associated with manual assessment.

IV. CONCLUSION

This study presented an Attention-Enhanced Deep Learning Framework (AEDL-RSS) for automating the rubber sheet grading using the VGG16 model. The Convolutional Block Attention Module (CBAM) was integrated by emphasizing spatial and channel-wise prominent features in the rubber sheet images, and validated using K-Fold cross-validation for 5 folds. The proposed method addresses challenges of subjectivity, inconsistency, and limited reproducibility in manual grading. This attention mechanism improved the network's ability to identify and capture fine-grained features crucial for grading. The 5-Fold cross-validation guaranteed the robust performance of the model, minimizing the bias with increased accuracy.

The experimental results validated substantial enhancements in overall test accuracy up to 95.05% compared to the traditional approaches, based on thresholding, simple neural networks, and LVQ classifiers [4-6], as well as earlier automated systems such as RSSGS [7]. The AEDL-RSS framework offers a proficient balance between fine-grained texture recognition, computational efficiency, and interpretability.

The novelty of this work lies in the integration of CBAM with a VGG16 architecture for Ribbed Smoked Rubber Sheets (RSS) grading, the introduction of an expert-annotated dataset, and the use of cross-validation for reliable performance, none of which appeared in prior RSS studies. The ability of the framework to classify the visually similar low-grade and mid-grade samples accurately substantiates its capability for real-world deployment in rubber processing units where consistent quality assessment is crucial.

The model demonstrates strong performance across most grades, and misclassification lies between visually similar grades, suggesting the need for more diverse real-world samples for further improvements. Future extensions may explore lightweight architectures for edge deployment and the integration of explainability tools for trust in automated grading decisions.

The proposed AEDL-RSS framework demonstrates robust generalization and provides a promising foundation for advancements in rubber sheet grading systems, which improve quality control, reduce human bias, and are suitable for real-time deployment in the rubber industries.

DATA AVAILABILITY STATEMENT

The dataset used in this study cannot be shared publicly due to proprietary restrictions. However, they are available from the corresponding author upon reasonable request.

REFERENCES

[1] N. Hütten, M. Alves Gomes, F. Hölken, K. Andricevic, R. Meyes, and T. Meisen, "Deep Learning for Automated Visual Inspection in Manufacturing and Maintenance: A Survey of Open-Access Papers," *Applied System Innovation*, vol. 7, no. 1, Jan. 2024, Art. no. 11, https://doi.org/10.3390/asi7010011.

[2] Z. Niu, G. Zhong, and H. Yu, "A Review on the Attention Mechanism of Deep Learning," *Neurocomputing*, vol. 452, pp. 48–62, Sep. 2021, https://doi.org/10.1016/j.neucom.2021.03.091.

[3] D. Soydaner, "Attention Mechanism in Neural Networks: Where It Comes and Where It Goes," 2022, https://doi.org/10.48550/ARXIV.2204.13154.

[4] I. Fibriani, Sumardi, A. B. Satriya, and S. B. Utomo, "Colour Based Image Processing Method for Recognizing Ribbed Smoked Sheet Grade," *IOP Conference Series: Materials Science and Engineering*, vol. 185, Mar. 2017, Art. no. 012025, https://doi.org/10.1088/1757-899X/185/1/012025.

[5] R. F. Rahmat, A. F. Pulungan, S. Faza, and R. Budiarto, "Image Classification of Ribbed Smoked Sheet Using Learning Vector Quantization," *Journal of Physics: Conference Series*, vol. 801, Jan. 2017, Art. no. 012050, https://doi.org/10.1088/1742-6596/801/1/012050.

[6] Z. Arifin, E. Y. Madya Mukti, A. Jehan Tamamy, and M. Ary Heryanto, "Classification of RSS (Ribbed Smoke Sheet) Based on Presence or Absence of Fungus Using the NN (Neural Network) Perceptron Method," in *2020 International Seminar on Application for Technology of Information and Communication,* Semarang, Indonesia, Sep. 2020, pp. 359–363, https://doi.org/10.1109/iSemantic50169.2020.9234277.

[7] C. Pornpanomchai and N. Chantharangsikul, "Ribbed Smoked Sheet Grading System (RSSGS)," in *2010 International Conference on Electronics and Information Engineering*, Kyoto, Japan, Aug. 2010, pp. V1-128-V1-132, https://doi.org/10.1109/ICEIE.2010.5559907.

[8] B. Suksawat and P. Komkum, "Pineapple Quality Grading Using Image Processing and Fuzzy Logic Based on Thai Agriculture Standards," in *2015 International Conference on Control, Automation and Robotics*, Singapore, May 2015, pp. 218–222, https://doi.org/10.1109/ICCAR.2015.7166035.

[9] H. S. Choi, J. B. Cho, S. G. Kim, and H. S. Choi, "A Real-Time Smart Fruit Quality Grading System Classifying by External Appearance and Internal Flavor Factors," in *2018 IEEE International Conference on Industrial Technology*, Lyon, France, Feb. 2018, pp. 2081–2086, https://doi.org/10.1109/ICIT.2018.8352510.

[10] H. Wang, "Garbage Recognition and Classification System Based on Convolutional Neural Network VGG16," in *2020 3rd International Conference on Advanced Electronic Materials, Computers and Software Engineering*, Shenzhen, China, Apr. 2020, pp. 252–255, https://doi.org/10.1109/AEMCSE50948.2020.00061.

[11] S. Choudhury and A. Agarwal, "Mango Disease Classification Using CNN," in *International Conference on Emerging Techniques in Computational Intelligence*, Hyderabad, India, pp. 1-6, Sep. 2023.

[12] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional Block Attention Module," in *Computer Vision – ECCV 2018*, vol. 11211, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham, Switzerland: Springer International Publishing, 2018, pp. 3–19.

[13] R. Kohavi, "A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection," in *14th International Joint Conference on Artificial Intelligence*, Montréal, QC, Canada, Aug. 1995, pp. 1137–1143.

[14] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition." arXiv, 2014, https://doi.org/10.48550/ARXIV.1409.1556.