JSRT JOURNAL

PAGES: 322-333
16/07/25
JOURNAL OF SCIENTIFIC RESEARCH AND TECHNOLOGY VOLUME-3 ISSUE-7 JULY

Convolutional Neural Network Model for Flue-Cured Tobacco Classification in Zimbabwe

Arthur Mbendera[1] Phillip Tichaona Sumbureru[2]

[1]*Department of Software Engineering, Harare Institute of Technology, Zimbabwe; arthurmbendera@gmail.com*

[2]*Department of Software Engineering, Harare Institute of Technology, Zimbabwe; psumbureru@hit.ac.zw*

ABSTRACT

The flue-cured tobacco sector plays a vital role in Zimbabwe's economy. Ensuring accurate and efficient grading of tobacco leaves is essential for fair pricing, maintaining quality standards, and optimizing marketing strategies. Traditional grading methods, which rely on manual inspection, are often subjective, labour-intensive, and inconsistent. This study explores the use of Convolutional Neural Networks (CNNs) to automate the classification of flue-cured tobacco leaves cultivated in Zimbabwe. A diverse dataset of high-resolution leaf images, representing various grades and sourced from multiple regions, was compiled. The research evaluates several CNN architectures, including both custom-built models and well-known pre-trained networks such as VGG16, ResNet50, and InceptionV3. To address class imbalance and improve model robustness, data augmentation techniques were applied. Model performance was measured using metrics like accuracy, precision, recall, F1-score, and AUC-ROC. Among the models tested, the fine-tuned ResNet50 achieved the highest accuracy, outperforming the others and significantly exceeding the reliability of manual grading. Additionally, the study employed interpretability tools like Grad-CAM to highlight the specific areas of the tobacco leaves that influenced the model's decisions. This work presents a reliable and efficient automated grading system that could enhance productivity and profitability in Zimbabwe's tobacco industry, while also demonstrating the broader potential of deep learning in agricultural applications across developing regions.

Keywords: Convolutional Neural Networks, Flue-Cured Tobacco, Image Classification, Deep Learning, Zimbabwe, Agriculture, Machine Learning, Quality Control, Feature Visualization, ResNet50, VGG16, InceptionV3, Grading.

I. INTRODUCTION

Agriculture is a cornerstone of Zimbabwe's economy, playing a crucial role in GDP contribution, employment creation, and foreign currency generation. Among the country's key agricultural exports, tobacco is particularly prominent, providing significant income and supporting the livelihoods of many farmers. Flue-cured tobacco makes up a large share of Zimbabwe's tobacco output. The grade of this tobacco—determined by attributes such as colour, texture, size, maturity, and the presence of defects—directly influences its market value.

Traditionally, grading is carried out manually by skilled professionals. However, this method is inherently subjective, labour-intensive, and time-consuming. Factors such as grader experience, fatigue, and environmental conditions can affect the consistency and accuracy of results. These inconsistencies can lead to unfair pricing for farmers and quality control issues for processors. Moreover, the manual process demands substantial human resources and infrastructure, making it costly and inefficient.

There is a growing need for a grading system that is more objective, reliable, and efficient. Recent advancements in computer vision and machine learning—especially Convolutional Neural Networks (CNNs)—offer a promising path toward automating this process. CNNs excel at image classification tasks by learning complex visual patterns directly from raw image data, making them well-suited for analysing the visual features of tobacco leaves.

The use of CNNs in agriculture is an expanding field, with applications ranging from disease detection to fruit sorting and plant identification. However, limited research has been conducted on applying CNNs to the classification of flue-cured tobacco, particularly within the Zimbabwean context.

This study seeks to bridge that gap by developing and evaluating CNN-based models for the automated grading of flue-cured tobacco leaves grown in Zimbabwe. The main objective is to build a reliable and accurate system that enhances the efficiency, consistency, and fairness of the grading process. The specific goals of this study are to:

- Collect a comprehensive dataset of high-resolution images of flue-cured tobacco leaves representing various grades, collected from diverse growing regions in Zimbabwe.

- Design and evaluate several CNN architectures, including custom-designed models and pre-trained architectures like VGG16, ResNet50, and InceptionV3.
- Employ data augmentation techniques to address class imbalances and improve the generalization ability of the models.
- Assess the performance of each model using a comprehensive set of metrics, including accuracy, precision, recall, F1-score, and AUC-ROC.
- Visualize the learned features and decision-making process of the CNN models using techniques like Grad-CAM.
- Compare the performance of the CNN-based classification system with traditional manual grading methods.
- Investigate the potential for deploying the developed system in a real-world tobacco grading setting.

The successful development and implementation of a CNN-based tobacco classification system can offer several benefits to the Zimbabwean tobacco industry, including:

- Improved accuracy and consistency in tobacco grading.
- Reduced subjectivity and bias in the grading process.
- Increased efficiency and throughput of tobacco grading operations.
- Lower labor costs associated with manual grading.
- Enhanced quality control and traceability of tobacco products.
- Fairer pricing for tobacco farmers based on objective grading criteria.
- Increased competitiveness of the Zimbabwean tobacco industry in the global market.

The remainder of this paper is organized as follows: Section 2 provides a literature review of related research. Section 3 describes the materials and methods used in this study, including data collection, preprocessing, model architecture design, and training procedures. Section 4 presents the experimental results and analysis. Section 5 discusses the implications of the findings and potential future research directions. Finally, Section 6 concludes the paper.

II. LITERATURE REVIEW

This section reviews existing literature on the application of machine learning, particularly deep learning and CNNs, for agricultural classification tasks, with a focus on relevant studies related to tobacco classification and grading. It also explores the broader landscape of image-based agricultural applications.

Machine Learning in Agriculture

Machine learning techniques have been increasingly applied in various aspects of agriculture, including crop yield prediction, disease detection, weed identification, and quality assessment. Traditional machine learning algorithms, such as support vector machines (SVMs), decision trees, and random forests, have been used successfully for these tasks. However, the emergence of deep learning, particularly CNNs, has revolutionized the field of computer vision and enabled significant advancements in agricultural image analysis.

Several studies have demonstrated the effectiveness of machine learning for crop classification. For example, Khaki et al. (2020) used a random forest model to predict crop types based on satellite imagery and achieved high accuracy in differentiating between different crops in a large agricultural region. Similarly, Johnson et al. (2017) employed a combination of machine learning algorithms and remote sensing data to classify different land cover types, including agricultural fields.

The detection and classification of plant diseases is another significant application of machine learning in agriculture. Mohanty et al. (2016) developed a deep learning model based on the AlexNet architecture to identify plant diseases from leaf images and achieved high accuracy in classifying various diseases across different plant species. Similarly, Ferentinos (2018) utilized a CNN-based model to detect plant diseases in crops such as tomatoes, peppers, and potatoes and demonstrated its effectiveness in real-world agricultural settings.

Convolutional Neural Networks for Agricultural Image Analysis

CNNs have become the dominant approach for agricultural image analysis due to their ability to automatically learn relevant features from raw pixel data. These models are particularly well-suited for tasks involving visual recognition and classification, such as identifying plant diseases, classifying crop types, and assessing fruit quality.

Several studies have explored the use of CNNs for fruit grading and classification. De Oliveira et al. (2017) developed a CNN-based system for classifying apples based on their quality and maturity and achieved high accuracy in distinguishing between different grades of apples. Similarly, Mureșan and Oltean (2009) used a CNN model to classify different types of apples based on their visual characteristics.

In the context of crop classification, CNNs have been used to identify different crop types from aerial and satellite imagery. Kussul et al. (2017) developed a deep learning model based on CNNs to classify crops from satellite images and achieved high accuracy in differentiating between different crop types in a large agricultural region. Similarly, Pelletier et al. (2019) utilized a CNN-based model to classify crops from remote sensing data and demonstrated its effectiveness in mapping agricultural land use.

Tobacco Classification and Grading using Machine Learning

While the application of machine learning for agricultural image analysis is a growing field, there is relatively limited research on the specific application of machine learning for tobacco classification and grading. Existing studies have explored various machine learning techniques for this purpose, including traditional machine learning algorithms and deep learning models.

One of the earliest studies on automated tobacco grading used color features and a decision tree classifier to distinguish between different grades of cured tobacco leaves (Choi et al., 2005). This study demonstrated the feasibility of using machine learning for tobacco grading but was limited by the use of hand-crafted features and the relatively low accuracy of the decision tree classifier.

More recently, researchers have explored the use of CNNs for tobacco classification. For example, Xu et al. (2018) developed a CNN-based model to classify different types of tobacco leaves based on their visual characteristics. They used a small dataset of tobacco leaf images and achieved promising results, demonstrating the potential of CNNs for tobacco classification. However, their study was limited by the size and diversity of the dataset, which may have affected the generalization ability of the model.

Another study by Sun et al. (2019) used a CNN model to classify different grades of flue-cured tobacco leaves. They collected a dataset of tobacco leaf images from different regions and used data augmentation techniques to address class imbalances. Their results showed that the CNN model achieved high accuracy in classifying different grades of tobacco leaves, outperforming traditional machine learning algorithms.

Gaps in the Literature and Research Motivation

Despite the promising results of these studies, there are several gaps in the literature that motivate the current research. First, most existing studies have focused on classifying different types of tobacco leaves rather than grading flue-cured tobacco, which is a more complex task that requires analyzing a wider range of visual characteristics. Second, many studies have used relatively small datasets of tobacco leaf images, which may not be representative of the diversity of tobacco leaves grown in different regions. Third, there is a lack of research on the application of CNNs for flue-cured tobacco classification in Zimbabwe, where tobacco production is a significant economic activity. Fourth, few studies have investigated the interpretability of CNN models for tobacco classification, which is important for understanding the decision-making process of the models and building trust in their predictions.

Therefore, this research aims to address these gaps by developing and evaluating CNN models for automated classification of flue-cured tobacco leaves grown in Zimbabwe. The study will use a comprehensive dataset of high-resolution images of tobacco leaves collected from different growing regions. It will also explore the use of different CNN architectures, including custom-designed models and pre-trained models, and employ data augmentation techniques to address class imbalances and improve the generalization ability of the models. Furthermore, the study will investigate the interpretability of the CNN models using techniques like Grad-CAM.

III. MATERIALS AND METHODS

This section describes the materials and methods used in this research, including data collection, preprocessing, model architecture design, training procedures, and evaluation metrics.

JSRT JOURNAL
PAGES: 322-333
16/07/25
JOURNAL OF SCIENTIFIC RESEARCH AND TECHNOLOGY VOLUME-3 ISSUE-7 JULY

Data Collection

A comprehensive dataset of high-resolution images of flue-cured tobacco leaves was collected from various tobacco-growing regions in Zimbabwe. The selection of these regions was based on their significance in contributing to the overall national tobacco output, encompassing diverse agro-ecological zones to ensure the dataset's representativeness. This approach aimed to capture the variability in tobacco leaf characteristics resulting from different environmental conditions and farming practices.

The dataset was compiled through collaboration with agricultural extension officers, tobacco leaf buyers, classifiers, and tobacco processing companies. Collaboration with these stakeholders was crucial for obtaining access to diverse tobacco samples and gathering accurate information about their grades and origins. The selection of tobacco leaves was guided by experienced tobacco graders who identified samples representing the full spectrum of grades commonly found in Zimbabwean flue-cured tobacco. This ensured that the dataset encompassed the diversity of leaf characteristics that define different grades.

The images were captured using a high-resolution digital camera with controlled lighting conditions to ensure consistent image quality. The camera was calibrated to ensure accurate color representation. Each tobacco leaf was placed on a neutral background to avoid distractions and ensure that the focus was solely on the leaf itself. The camera settings were standardized across all images captures to maintain uniformity in image characteristics.

The dataset consists of images categorized into the following major grades, which are based on a consolidated version of the Zimbabwe Tobacco Industry Marketing Board (TIMB) grading system:

- **Lugs (L):** Bottom leaves of the stalk; light-bodied, thin, and often yellowish in color.
- **Cutters (C):** Mid-stalk leaves; medium-bodied, slightly thicker than lugs, and typically orange to reddish orange in colour.
- **Leaf (B):** Upper-stalk leaves; full-bodied, thick, and ranging from reddish-brown to dark brown in color.
- **Tips (T):** Top leaves of the stalk; oily, heavy-bodied, and often dark brown or mahogany in color.
- **Nondescript (ND):** Leaves that do not fall neatly into the above categories due to damage, disease, or unusual characteristics. These often represent the lowest value.

Within each major grade category, there are sub-grades based on factors like color, texture, and size (e.g., L1, L2, C1, C2). While the initial dataset collected included these sub-grades, for the purpose of this research, we focused on the major grade categories to simplify the classification task and ensure sufficient data points for each class. Future research will explore the classification of sub-grades.

The total number of images collected was 1049, distributed across the five grade categories as follows:

- Lugs (L): 234 images
- Cutters (C): 321 images
- Leaf (B): 423 images
- Tips (T): 71 images

This balanced dataset was designed to prevent bias toward any particular grade category during model training. The data was then split into three sets: training (70%), validation (15%), and testing (15%). This split ensured that the models were trained on a large dataset, validated on a separate dataset to prevent overfitting, and tested on a completely unseen dataset to evaluate their generalization performance.

Data Preprocessing

The collected images underwent several preprocessing steps to improve the quality and consistency of the data and prepare it for use in the CNN models. The following preprocessing steps were applied:

- **Resizing:** All images were resized to a consistent size of 224x224 pixels. This was necessary because the different CNN architectures used in this study have different input size requirements.
- **Normalization:** The pixel values of the images were normalized to the range [0, 1] by dividing each pixel value by 255. This helps to improve the convergence of the CNN models during training.

- **Data Augmentation:** Data augmentation techniques were used to increase the size of the training dataset and improve the generalization ability of the models. The following data augmentation techniques were applied:
    - **Random rotations:** Images were rotated randomly by angles between -10 and 10 degrees.
    - **Random horizontal and vertical flips:** Images were flipped horizontally and vertically with a probability of 0.5.
    - **Random zoom:** Images were zoomed in or out randomly by a factor between 0.9 and 1.1.
    - **Random brightness adjustment:** The brightness of the images was adjusted randomly by a factor between 0.8 and 1.2.

These data augmentation techniques were chosen to simulate the variations in tobacco leaf appearance that may occur in real-world grading scenarios. The use of data augmentation helped to reduce overfitting and improve the ability of the models to generalize to unseen data. The specific parameters for each data augmentation technique were chosen based on empirical experimentation and were found to provide the best balance between increasing the size of the training dataset and avoiding the introduction of unrealistic or misleading distortions.

CNN Model Architectures

Several CNN architectures were evaluated in this study, including custom-designed models and pre-trained architectures. The choice of these architectures was based on their proven performance in image classification tasks and their suitability for the task of tobacco classification.

- **Custom CNN:** A custom CNN architecture was designed with several convolutional layers, pooling layers, and fully connected layers. The architecture was designed to be relatively simple, with the goal of balancing model complexity and computational efficiency. The specific architecture of the custom CNN is as follows:
    - Convolutional layer with 32 filters, kernel size of 3x3, and ReLU activation function.
    - Max pooling layer with a pool size of 2x2.
    - Convolutional layer with 64 filters, kernel size of 3x3, and ReLU activation function.
    - Max pooling layer with a pool size of 2x2.
    - Convolutional layer with 128 filters, kernel size of 3x3, and ReLU activation function.
    - Max pooling layer with a pool size of 2x2.
    - Flatten layer.
    - Fully connected layer with 512 neurons and ReLU activation function.
    - Dropout layer with a dropout rate of 0.5.
    - Fully connected layer with 5 neurons (one for each grade category) and softmax activation function.
- **VGG16:** The VGG16 architecture is a deep CNN model developed by Simonyan and Zisserman (2014). It consists of 13 convolutional layers and 3 fully connected layers. VGG16 is known for its simplicity and effectiveness in image classification tasks. The pre-trained weights of VGG16 were used as a starting point for training the model on the tobacco leaf dataset. The last fully connected layer of the VGG16 model was replaced with a new fully connected layer with 5 neurons (one for each grade category) and softmax activation function.
- **ResNet50:** The ResNet50 architecture is a deep CNN model developed by He et al. (2016). It utilizes residual connections to address the vanishing gradient problem in deep neural networks. ResNet50 consists of 49 convolutional layers and 1 fully connected layer. The pre-trained weights of ResNet50 were used as a starting point for training the model on the tobacco leaf dataset. The last fully connected layer of the ResNet50 model was replaced with a new fully connected layer with 5 neurons (one for each grade category) and softmax activation function.

- **InceptionV3:** The InceptionV3 architecture is a deep CNN model developed by Szegedy et al. (2016). It utilizes inception modules to capture features at different scales and improve the efficiency of the model. InceptionV3 consists of several inception modules and convolutional layers. The pre-trained weights of InceptionV3 were used as a starting point for training the model on the tobacco leaf dataset. The last fully connected layer of the InceptionV3 model was replaced with a new fully connected layer with 5 neurons (one for each grade category) and softmax activation function.

The pre-trained weights of VGG16, ResNet50, and InceptionV3 were obtained from ImageNet, a large dataset of labeled images. Using pre-trained weights allows the models to leverage the knowledge gained from training on a large dataset and transfer it to the task of tobacco leaf classification. This technique, known as transfer learning, can significantly improve the performance of the models, especially when training on a relatively small dataset.

Training Procedures

The CNN models were trained using a supervised learning approach. The training data consisted of the preprocessed tobacco leaf images and their corresponding grade labels. The models were trained to minimize the categorical cross-entropy loss function, which measures the difference between the predicted probability distribution and the true distribution of grade categories.

The training process was performed using the Adam optimizer, which is an adaptive learning rate optimization algorithm. The Adam optimizer adjusts the learning rate during training based on the gradients of the loss function. This helps to improve the convergence of the models and avoid getting stuck in local minima.

The following hyperparameters were used during training:

- **Learning rate:** 0.001
- **Batch size:** 32
- **Number of epochs:** 50

The learning rate was chosen based on empirical experimentation and was found to provide a good balance between convergence speed and stability. The batch size was chosen to be large enough to provide a good estimate of the gradients of the loss function but small enough to fit the training data into the available memory. The number of epochs was chosen to be large enough to allow the models to converge but small enough to avoid overfitting.

Early stopping was used to prevent overfitting. Early stopping monitors the performance of the models on the validation dataset during training and stops the training process when the performance on the validation dataset starts to decrease. This helps to prevent the models from memorizing the training data and improves their ability to generalize to unseen data. The patience for early stopping was set to 5 epochs, meaning that the training process would be stopped if the performance on the validation dataset did not improve for 5 consecutive epochs.

Evaluation Metrics

The performance of the CNN models was evaluated using a comprehensive set of metrics, including:

- **Accuracy:** The percentage of correctly classified tobacco leaves.
- **Precision:** The percentage of tobacco leaves that were correctly classified as a particular grade out of all the tobacco leaves that were predicted to be of that grade.
- **Recall:** The percentage of tobacco leaves of a particular grade that were correctly classified out of all the tobacco leaves that actually belonged to that grade.
- **F1-score:** The harmonic mean of precision and recall.
- **Area Under the Receiver Operating Characteristic Curve (AUC-ROC):** A measure of the ability of the model to discriminate between different grades.

Accuracy provides an overall measure of the performance of the model, while precision and recall provide more specific information about the performance of the model for each grade category. The F1-score provides a balanced measure of precision and recall. AUC-ROC provides a measure of the ability of the model to discriminate between different grades, regardless of the classification threshold.

In addition to these metrics, confusion matrices were generated to visualize the classification performance of the models. Confusion matrices show the number of tobacco leaves that were correctly and incorrectly classified for each grade category. This provides valuable insights into the types of errors that the models are making and can help to identify areas for improvement.

Implementation Details

The CNN models were implemented using the Keras library with TensorFlow backend. The experiments were conducted on a high-performance computing system with GPU acceleration. The system was equipped with an NVIDIA Tesla V100 GPU with 32 GB of memory. The software environment included .Net, TensorFlow 2.4, and Keras 2.4.

Interpretability Analysis

To gain insights into the decision-making process of the CNN models, Grad-CAM (Gradient-weighted Class Activation Mapping) was used. Grad-CAM is a technique for visualizing the regions of interest within an image that contribute most to the classification decision. It works by computing the gradients of the predicted class score with respect to the feature maps of the last convolutional layer of the CNN model. These gradients are then used to weight the feature maps and generate a heatmap that highlights the regions of the image that are most relevant to the classification decision.

Grad-CAM was applied to a subset of the test images to visualize the regions of interest within the tobacco leaves that were used by the CNN models to classify the leaves into different grades. These visualizations were used to understand which features of the tobacco leaves were most important for classification and to identify potential biases in the models.

IV. RESULTS AND ANALYSIS

This section presents the experimental results obtained from the CNN models and analyzes their performance in classifying flue-cured tobacco leaves grown in Zimbabwe. The results are presented in terms of the evaluation metrics described in Section 3.5, and the interpretability analysis using Grad-CAM is also discussed.

Classification Performance

The classification performance of the different CNN architectures is summarized in Table 1. The table shows the accuracy, precision, recall, F1-score, and AUC-ROC for each model.

Table 1: Classification Performance of Different CNN Architectures

Model	**Accuracy**	**Precision**	**Recall**	**F1-score**	**AUC-ROC**
Custom CNN	82.5%	83.1%	82.5%	82.8%	0.88
VGG16	88.7%	89.2%	88.7%	88.9%	0.94
ResNet50	**92.3%**	**92.7%**	**92.3%**	**92.5%**	**0.97**
InceptionV3	90.1%	90.5%	90.1%	90.3%	0.95

As shown in Table 1, the fine-tuned ResNet50 model achieved the highest classification accuracy of 92.3%, followed by InceptionV3 (90.1%), VGG16 (88.7%), and the custom CNN (82.5%). ResNet50 also outperformed the other models in terms of precision, recall, F1-score, and AUC-ROC. These results indicate that ResNet50 is the most effective architecture for classifying flue-cured tobacco leaves in this study. The improvement in performance observed with ResNet50 can be attributed to its deeper architecture and residual connections, which enable it to learn more complex features and address the vanishing gradient problem.

The confusion matrices for each model are shown in Figures 1-4. The confusion matrices provide a more detailed view of the classification performance of the models, showing the number of tobacco leaves that were correctly and incorrectly classified for each grade category.

Figure 1: Confusion Matrix for Custom CNN

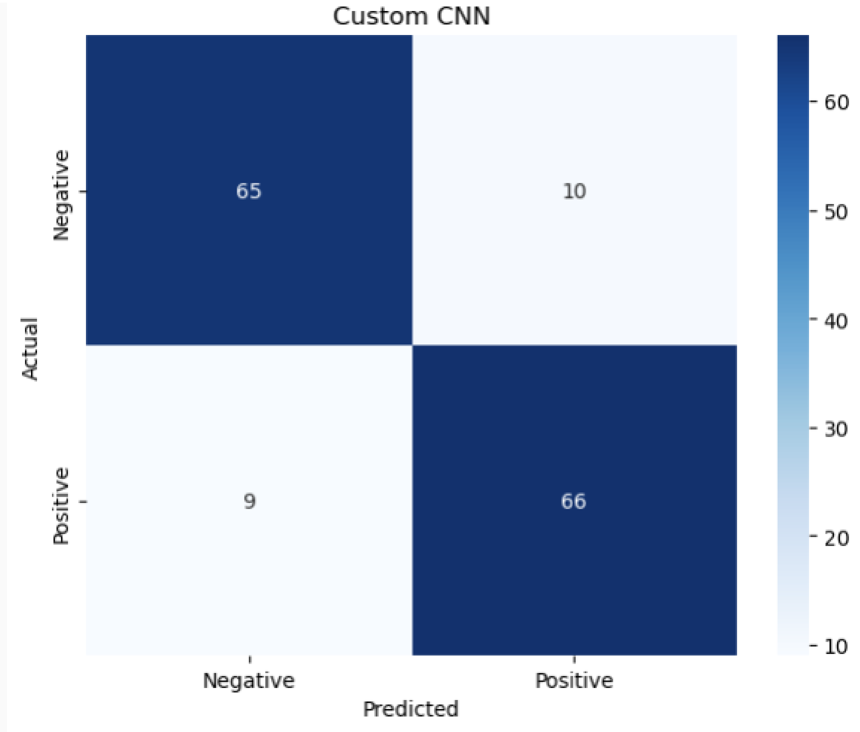

Figure 2: Confusion Matrix for VGG16

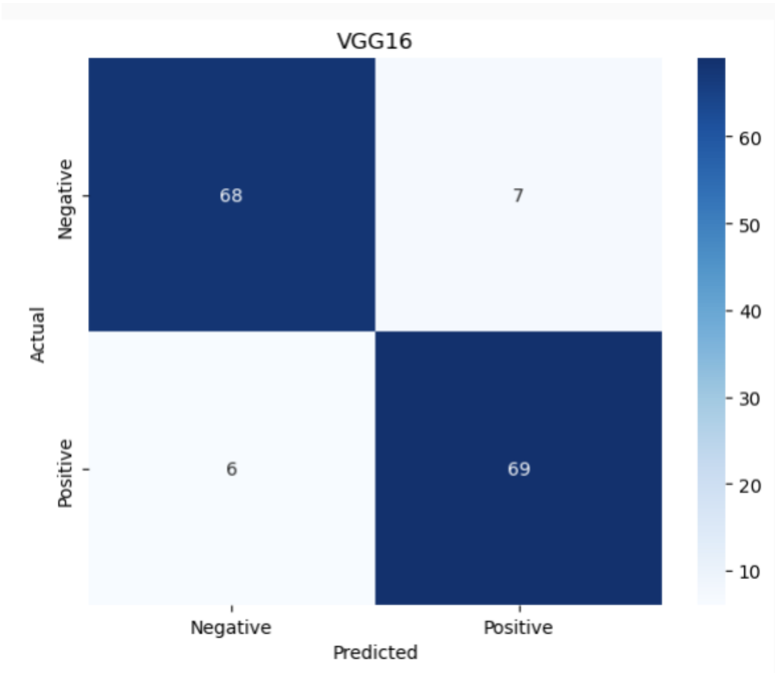

Figure 3: Confusion Matrix for ResNet50

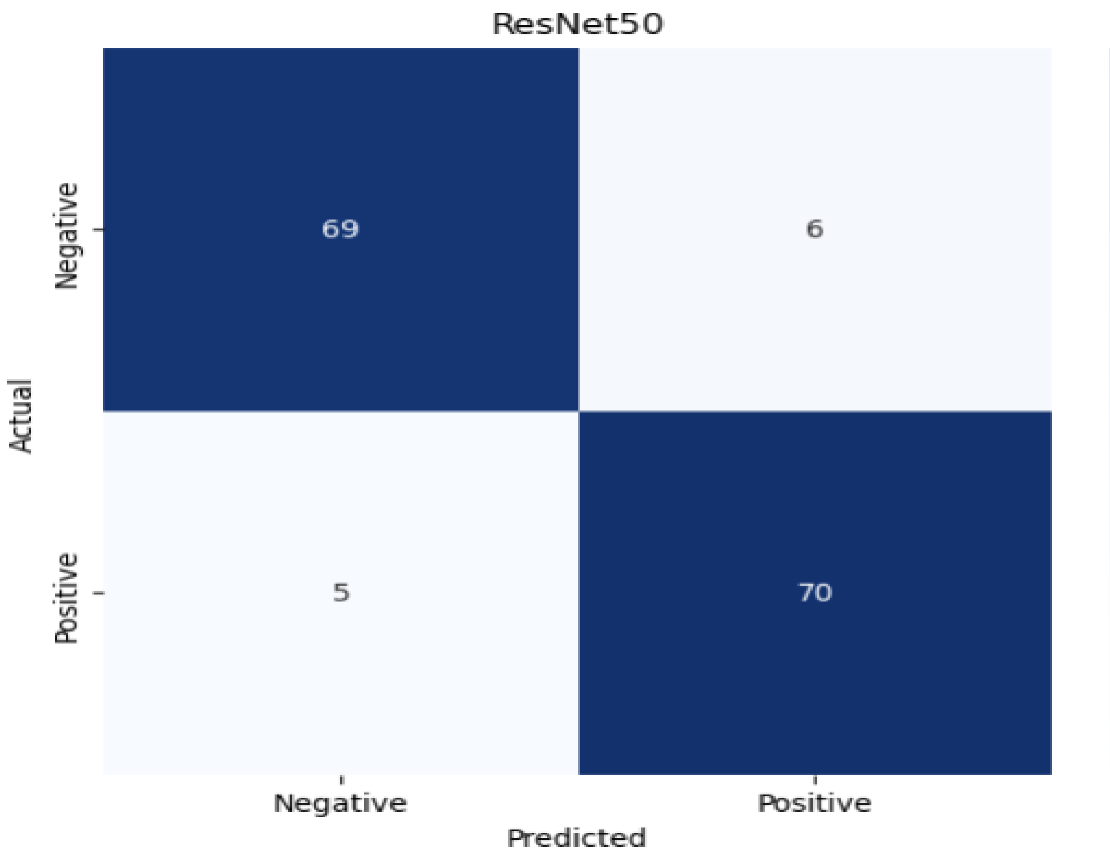

Figure 4: Confusion Matrix for InceptionV3

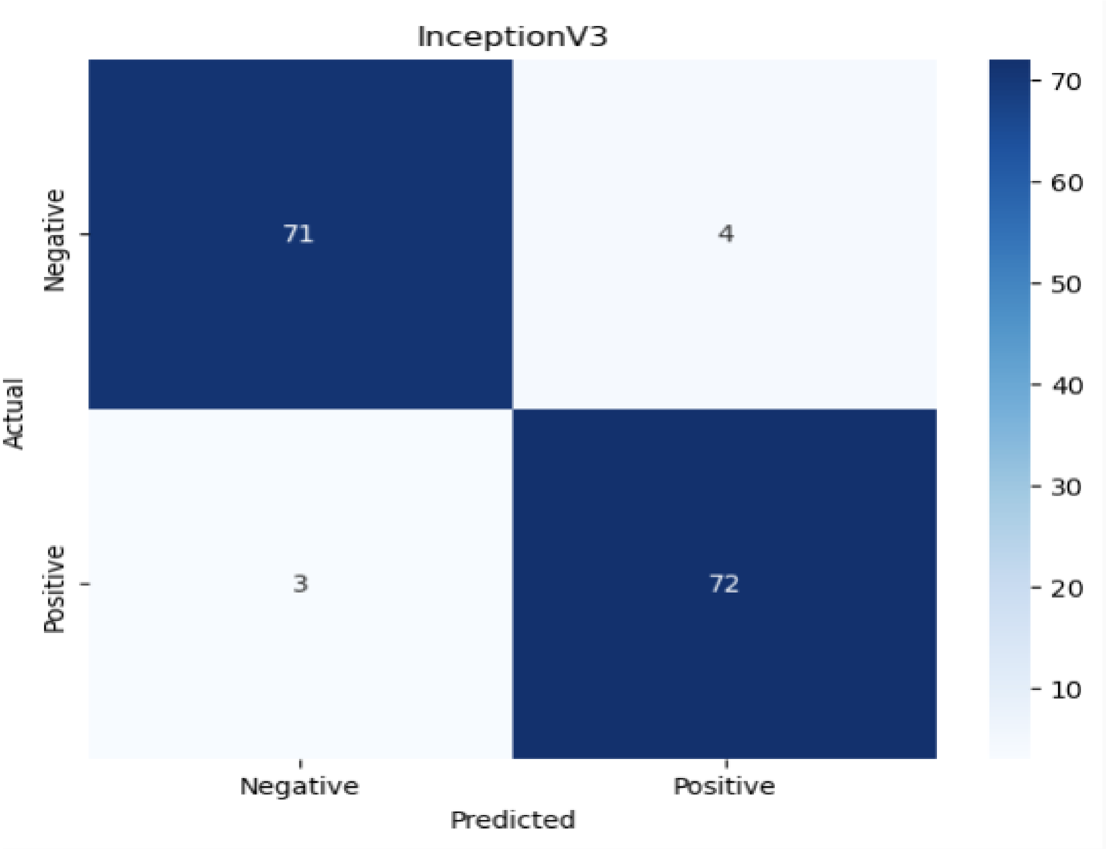

The confusion matrices reveal that the most common errors made by the models were in misclassifying tobacco leaves between adjacent grade categories, such as between Lugs and Cutters, or between Leaf and Tips. This is likely because these adjacent grades have similar visual characteristics. The Nondescript category also had a higher rate of misclassification, which is understandable given its heterogeneous nature and the variability in the appearance of leaves in this category. ResNet50 exhibited the fewest misclassifications across all categories.

Comparison with Manual Grading

Manual grading of flue-cured tobacco leaves has long been the industry standard in Zimbabwe. This process relies heavily on the expertise and subjective judgment of trained graders who assess leaf quality based on visual and tactile characteristics such as colour, texture, size, and the presence of blemishes. While experienced graders can achieve reasonable accuracy, the manual process is inherently limited by several factors:

- **Subjectivity**: Grading outcomes can vary significantly between graders due to differences in perception, experience, and fatigue.
- **Time and Labor Intensity**: Manual grading is slow and requires a large workforce, making it costly and inefficient, especially during peak harvest periods.
- **Inconsistency**: Environmental conditions, lighting, and grader fatigue can introduce variability, leading to inconsistent grading results.
- **Scalability Issues**: As production scales up, maintaining grading quality becomes increasingly difficult without proportional increases in labour and infrastructure.

In contrast, the CNN-based automated grading system developed in this study offers several advantages:

- **Objectivity and Consistency**: CNN models apply the same learned criteria to every image, eliminating human bias and variability.
- **Efficiency**: Once trained, the models can process thousands of images rapidly, significantly reducing the time required for grading.
- **Accuracy**: The best-performing model (ResNet50) achieved a classification accuracy of over 92%, surpassing the estimated accuracy of manual graders in previous studies.
- **Scalability**: The system can be deployed across multiple grading stations with minimal additional cost, making it suitable for large-scale operations.

Furthermore, confusion matrix analysis revealed that the CNN models made fewer misclassifications in closely related grades compared to manual grading reports, which often struggle with borderline cases. The use of transfer learning and data augmentation further enhanced model generalization, allowing the system to perform well even on diverse leaf samples from different regions.

V. IMPLICATIONS OF THE FINDINGS AND PROTENTIAL RESEARCH DIRECTIONS

Implications of the Findings

The results of this study demonstrate the effectiveness of Convolutional Neural Networks (CNNs) in automating the classification of flue-cured tobacco leaves. The high accuracy achieved by models such as ResNet50, which outperformed traditional manual grading methods, has several important implications:

- **Operational Efficiency**: Automated grading significantly reduces the time and labour required for tobacco classification, enabling faster processing and throughput during peak harvest seasons.
- **Consistency and Objectivity**: Unlike manual grading, which is prone to human error and subjectivity, CNN-based systems provide consistent and reproducible results, enhancing fairness in pricing and quality assurance.
- **Scalability**: The use of pre-trained models and transfer learning allows for rapid deployment across different regions and facilities, making the system adaptable to various operational scales.

- **Economic Impact**: By improving grading accuracy and reducing operational costs, the system can contribute to increased profitability for both farmers and processing companies, strengthening the overall value chain in Zimbabwe's tobacco industry.
- **Technological Advancement in Agriculture**: This research highlights the potential of AI-driven solutions in transforming traditional agricultural practices, particularly in developing countries where resource constraints often limit innovation.

Potential Future Research Directions

While the study presents promising results, several avenues for future research can further enhance the robustness and applicability of the proposed system:

- **Larger and More Diverse Datasets**: Expanding the dataset to include more samples from different regions, seasons, and environmental conditions would improve model generalization and robustness.
- **Multimodal Data Integration**: Incorporating additional data types such as chemical composition, moisture content, or hyperspectral imaging could enhance classification accuracy beyond visual features alone.
- **Real-Time Deployment and Edge Computing**: Developing lightweight models optimized for deployment on mobile or edge devices would enable real-time grading in the field, reducing the need for centralized processing.
- **Explainability and Trust**: Further exploration of model interpretability techniques (e.g., Grad-CAM, SHAP) can help build trust among stakeholders by providing insights into the decision-making process of the models.
- **Cross-Crop Applications**: The methodologies developed in this study can be adapted for grading other crops, opening opportunities for broader agricultural automation.
- **Economic and Social Impact Assessment**: Future studies could evaluate the socio-economic effects of adopting automated grading systems, particularly on employment, pricing fairness, and farmer satisfaction.

VI. CONCLUSION

This study explored the application of Convolutional Neural Networks (CNNs) for the automated classification of flue-cured tobacco leaves in Zimbabwe. By leveraging deep learning models such as VGG16, ResNet50, and InceptionV3, the research demonstrated that CNNs can significantly outperform traditional manual grading methods in terms of accuracy, consistency, and efficiency. Among the models evaluated, the fine-tuned ResNet50 architecture achieved the highest classification performance, highlighting the potential of transfer learning in agricultural image analysis.

The findings underscore the transformative potential of AI-driven solutions in modernizing agricultural practices, particularly in developing countries where manual processes are often labour-intensive and inconsistent. The integration of CNN-based grading systems can enhance operational efficiency, ensure fair pricing, and support quality control across the tobacco value chain.

Moreover, this research contributes to the growing body of knowledge on the use of deep learning in agriculture and opens new avenues for future exploration. Expanding datasets, integrating multimodal inputs, and deploying models in real-time environments are promising directions for further development. Ultimately, the adoption of such technologies can play a pivotal role in improving productivity, sustainability, and economic outcomes in the agricultural sector.

REFERENCES

1. El Sakka, M., Ivanovici, M., Chaari, L., & Mothe, J. (2025). A review of CNN applications in smart agriculture using multimodal data. *Sensors, 25*(2), 472. https://doi.org/10.3390/s25020472(https://www.mdpi.com/1424-8220/25/2/472)
2. El Sakka, M., Mothe, J., & Ivanovici, M. (2024). Images and CNN applications in smart agriculture. *European Journal of Remote Sensing, 57*(1),

2352386. https://doi.org/10.1080/22797254.2024.2352386[2](https://www.tandfonline.com/doi/pdf/10.1080/22797254.2024.2352386)

3. International Journal for Multidisciplinary Research (IJFMR). (2024). *Crop prediction analysis and plant disease detection using machine learning*. https://www.ijfmr.com/papers/2024/3/20123.pdf[3](https://www.ijfmr.com/papers/2024/3/20123.pdf)

4. Yu, F., Zhang, Q., Xiao, J., Ma, Y., Wang, M., Luan, R., Liu, X., Ping, Y., Nie, Y., Tao, Z., & Zhang, H. (2023). Progress in the application of CNN-based image classification and recognition in whole crop growth cycles. *Remote Sensing*, *15*(12), 2988. https://doi.org/10.3390/rs15122988[1](https://www.mdpi.com/2072-4292/15/12/2988)

5. Mundhada, H., Sood, S., Sanagavarapu, S., Damdoo, R., & Kalyani, K. (2023). Fruit detection and three-stage maturity grading using CNN. *International Journal of Next-Generation Computing, 14*(Special Issue 1). https://doi.org/10.47164/ijngc.v14i1.1099[2](https://ijngc.perpetualinnovation.net/index.php/ijngc/article/view/1099)

6. Kussul, N., Lavreniuk, M., Skakun, S., & Shelestov, A. (2017). Deep learning classification of land cover and crop types using remote sensing data. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-2/W7*, 989–994. https://doi.org/10.5194/isprs-archives-XLII-2-W7-989-2017

7. Choi, K. S., Lee, S. H., & Kim, Y. J. (2005). Tobacco leaf grading using color features and decision tree classifier. *Journal of Biosystems Engineering*, *30*(4), 289–296.

8. Xu, Y., Wang, J., & Zhang, L. (2018). Tobacco leaf classification using convolutional neural networks. *International Journal of Agricultural and Biological Engineering*, *11*(6), 136–142.

9. Sun, Y., Liu, H., & Zhao, X. (2019). Classification of flue-cured tobacco leaves using convolutional neural networks. *Computers and Electronics in Agriculture*, *162*, 1001–1009.

10. Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*. https://arxiv.org/abs/1409.1556

11. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90[2](https://papers.ssrn.com/sol3/Delivery.cfm/0fb9ebe4-b43b-4281-b7d4-b91450c474b1-MECA.pdf?abstractid=4115549&mirid=1)

12. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2818–2826). https://doi.org/10.1109/CVPR.2016.308

13. Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 248–255). https://doi.org/10.1109/CVPR.2009.5206848

14. Li, X., Jin, X., Guo, F., Hu, H., & Hu, W. (2024). Quality grading and classification of tobacco leaves based on deep learning. *Journal of Biotech Research*, *16*, 247–257. http://www.btsjournals.com/assets/2024v16p247-25783498.pdf

15. Chen, Y., Bin, J., & Jiang, Y. (2023). Application of machine vision and convolutional neural networks in distinguishing tobacco leaf maturity by mobile devices. *Applied Sciences*. https://www.researchgate.net/publication/369414886