International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4676-4683
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0614

An Automated Boiled Arecanut Grading Using Deep Learning

Keerthana K[1], *H L Shilpa*[2], *Kruthika M S*[3]
[1,3]*Department of MCA, PES College of Engineering, Mandya, Karnataka, India.*
[2]*Assistant Professor, Department of MCA, PES College of Engineering, Mandya, Karnataka, India.*
Email ID: *keerthanabhanumathi5@gmail.com*[1], *shilpa@pesce.ac.in*[2], *kruthikasuresh.m.p@gmail.com*[3]

Abstract
The commercial value of boiled arecanut is determined largely by its visual grade, and the process by which this grade is assigned continues, in most processing units, to depend on the manual judgement of experienced workers. Such an approach is slow, requires a substantial workforce, and is prone to inconsistency because human assessment varies with fatigue, lighting conditions, and individual perception. This paper presents an automated grading framework that applies deep learning and computer vision to classify boiled arecanut samples into four commercially recognized categories — Hasa, Bette, Rashidi, and Gorabalu — directly from digital images. A dataset of 1,792 images was collected under controlled illumination in Shivamogga district, Karnataka, and subjected to resizing, ImageNet-based normalization, and augmentation before being used to train three transfer-learning-based convolutional neural networks, namely AlexNet, GoogLeNet, and ResNet50. The three architectures were benchmarked against one another using accuracy, precision, recall, and F1-score, with AlexNet emerging as the strongest performer at a test accuracy of 99.15 percent. To make the system usable outside a research setting, a Streamlit-based web interface was built that allows a processing unit to upload an arecanut image and receive an immediate grade prediction along with a confidence score. Gradient-weighted Class Activation Mapping (Grad-CAM) was further incorporated to visualize the image regions that most influenced each prediction, giving operators a transparent view of the model's reasoning rather than a black-box output. Collectively, the results indicate that the proposed system is capable of replacing subjective manual grading with a fast, consistent, and interpretable alternative suited to real-world agricultural processing.

Keywords: *Boiled Arecanut Grading, Deep Learning, Computer Vision, Convolutional Neural Network (CNN), Transfer Learning, AlexNet, GoogLeNet, ResNet50, Streamlit, Grad-CAM, Explainable Artificial Intelligence (XAI), Agricultural Automation.*

1. Introduction

Agriculture continues to be the economic backbone of rural India, and the income of millions of farming households depends directly on the quality and market value of the crops they cultivate. Among the plantation crops grown across the southern and northeastern states, arecanut (Areca catechu L.), commonly referred to as betel nut, occupies a particularly important commercial position. Karnataka, Kerala, and Assam together account for the bulk of national production, and within these states the price a farmer or trader receives is tied closely to the grade assigned to the harvested nuts. Attributes such as colour uniformity, surface texture, size, and shape determine not only the price paid within the domestic market but also the crop's suitability for export and industrial processing, which is why consistent and trustworthy grading practices matter as much to the wider supply chain as they do to individual growers. In most processing centres, grading is still carried out by hand. Skilled workers inspect each nut and sort it into a commercial category based on what they see and, to some extent, feel. This method draws on years of accumulated experience, but it is far from efficient: it requires a large number of workers, particularly during the harvest season when volumes peak, and it inevitably introduces delays into the post-harvest workflow. More importantly, the outcome of a manual grading exercise is subjective. Two workers inspecting the same batch, or even the same worker at different

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4676-4683
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0614

times of day, can arrive at different conclusions depending on fatigue, ambient lighting, and personal judgement. The result is inconsistent batches, which in turn can lead to disputes over pricing and dissatisfaction among buyers. As processing volumes continue to grow, the case for a faster, more consistent, and less labour-dependent grading method becomes increasingly compelling. Advances in machine learning and computer vision have opened a practical route toward automating this task. Early work in the arecanut domain combined conventional image-processing pipelines with handcrafted feature extraction to separate nuts into quality categories [1]. Later studies showed that pretrained convolutional neural networks, adapted through transfer learning, could raise classification accuracy substantially compared with these handcrafted approaches [2]. Building on this, researchers introduced real-time segregation systems built around object detection models to speed up sorting and reduce the reliance on manual labour on the processing line [3], while other groups applied deeper architectures such as InceptionV3, EfficientNet, and ResNet to further improve grading accuracy [4]. Most recently, end-to-end platforms that combine intelligent vision sensing with automated sorting mechanisms have begun to appear in industrial settings [5]. Taken together, this body of work illustrates how far automated grading has progressed, and it also reveals where the remaining gaps lie. The wider progress of deep learning has been central to this shift. Modern CNN architectures can pick out subtle visual regularities in an image that are easy for a human inspector to miss or to judge inconsistently, and transfer learning allows a network already trained on a very large, general-purpose image dataset to be adapted to a specialized agricultural task using comparatively little domain-specific data. This drastically reduces both the data-collection burden and the computational cost of building a working classifier, which is one of the reasons deep learning has been adopted so widely across agricultural quality-assessment problems in recent years. Even so, a closer look at the existing literature shows that almost all of this research has been carried out on raw, dehusked, or Chali arecanut varieties. Boiled arecanut, which undergoes a distinct processing step that changes its surface appearance, colour, and texture, has received comparatively little attention, despite being commercially significant in its own right. In addition, few studies have gone beyond a single chosen architecture to benchmark multiple CNN models against one another on the same task, and fewer still have paired their classifiers with explainability tools that would let a user understand why a particular grade was assigned rather than simply accepting the output at face value. This paper addresses both gaps directly. It proposes a deep-learning-based grading system for boiled arecanut that classifies samples into four commercially recognized grades — Hasa, Bette, Rashidi, and Gorabalu — and it systematically compares three transfer-learning-based CNN architectures, AlexNet, GoogLeNet, and ResNet50, on this task. The resulting model is deployed through a Streamlit web application for real-time use, and its predictions are made interpretable through Gradient-weighted Class Activation Mapping (Grad-CAM), so that the regions of the image driving each decision are visible to the end user.

2. Literature Review

A range of studies has examined how machine learning and deep learning can be applied to arecanut grading, and reviewing them side by side helps to place the present work in context. One early framework combined image enhancement and segmentation with handcrafted descriptors — colour histograms, Gray Level Co-occurrence Matrix (GLCM) features, Local Binary Patterns (LBP), and shape descriptors — and reported a classification accuracy of 98.5 percent using a Gradient Boosting classifier [1]. A separate investigation moved toward deep feature learning by fine-tuning several pretrained CNNs, including InceptionV3, ResNet50, Inception-ResNetV2, and VGG16, for arecanut classification, finding that InceptionV3 gave the best testing accuracy at 91.98 percent [2]. Between them, these two studies demonstrate that both conventional feature engineering and modern transfer learning can be made to work reasonably well on arecanut imagery, even though their absolute performance differs. Other researchers have focused less on classification accuracy alone and more on getting a working system onto the processing line. A

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4676-4683
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0614

YOLOv5-based segregation pipeline deployed on a Raspberry Pi enabled on-the-fly detection and physical sorting of arecanuts, and reported competitive precision and recall under real operating conditions [3]. In a related study, a CNN ensemble built from InceptionV3, EfficientNetB3, and ResNet50 reached roughly 94 percent accuracy specifically for Chali arecanut grading, reinforcing the case for deep architectures when the visual differences between grades are subtle [4]. Industrial-scale efforts have pushed accuracy even higher: one CNN-based grading and sorting system achieved close to 98.23 percent accuracy while simultaneously automating the physical sorting step [5], and a separate texture-based approach relying on Wavelet features, Gabor filters, LBP, and GLCM descriptors reached 91.43 percent, showing that carefully engineered texture features still carry useful discriminative signal even without deep feature learning [6]. Looking beyond arecanut, similar CNN-based grading strategies have proven effective across a variety of other crops, which lends further confidence to the general approach adopted here. A maturity-grading framework for dragon fruit built on ResNet and MobileNet architectures achieved strong classification performance, illustrating how well these networks transfer across different crop types [7]. In the citrus domain, a system combining EfficientNet and VGG16 with image augmentation reported a sorting accuracy of about 99.73 percent, underscoring how much data augmentation can help when the available training data is limited [8]. A tomato-grading framework using EfficientNet-B3 and ResNet50 reached 99.03 percent accuracy while remaining light enough for real-time inference [9], and a mango-quality system that combined AlexNet, VGG16, and ResNet with saliency-map-based explainability techniques showed that interpretability and accuracy are not mutually exclusive goals [10]. Read together, these studies confirm that transfer-learning-based CNNs are a mature and reliable tool for agricultural grading, but they also expose a consistent blind spot in the arecanut literature specifically. Nearly every prior study has worked with raw, dehusked, or Chali varieties, leaving boiled arecanut — which is visually distinct on account of the boiling process — largely unaddressed. Furthermore, systematic side-by-side comparison of multiple CNN architectures on a single boiled-arecanut dataset, combined with an explainability mechanism such as Grad-CAM, has not previously been reported. The present work is designed to close precisely this gap, by building a transfer-learning-based classifier for boiled arecanut grading, benchmarking AlexNet, GoogLeNet, and ResNet50 against one another under identical conditions, and pairing the resulting model with Grad-CAM visualization to keep its decisions transparent.

3. Methodology

The proposed system follows a structured pipeline that begins with image acquisition and preprocessing, proceeds through dataset preparation and transfer-learning-based model development, and concludes with performance evaluation, web-based deployment, and explainable-AI visualization. Each of these stages is described in detail in the subsections that follow.

3.1. Dataset Description

The dataset used in this study consists of 1,792 high-resolution images of boiled arecanut samples, organized into four categories corresponding to the commercial grades Hasa, Bette, Rashidi, and Gorabalu. Each image captures the visual characteristics — colour, surface texture, and shape — that a human grader would ordinarily rely on, making the dataset well suited to a CNN-based classification approach [11- 15].

3.2. Image Acquisition

All images were captured through a structured data-collection exercise conducted in Shivamogga district, Karnataka, using a calibrated high-resolution digital camera. To keep the visual characteristics within each class as consistent as possible and to limit the influence of external factors on the images, every sample was photographed under uniform lighting against a plain, neutral-coloured background. This setup was chosen specifically to suppress shadow formation and surface reflections, both of which can otherwise distort the colour and texture cues that distinguish one grade from another.

3.3. Image Preprocessing

Before being used for training, every image in the dataset was passed through a preprocessing pipeline intended to improve data quality and to make the

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4676-4683
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0614

images compatible with the chosen CNN architectures. Each image was resized to a fixed resolution of 224 × 224 pixels and converted to a standardized format, and pixel intensities were normalized using ImageNet statistics so that the feature distribution seen by the network during fine-tuning matched what it had encountered during pretraining. To help the models generalize beyond the exact conditions under which the training images were captured, augmentation operations — random rotation, horizontal flipping, zooming, and cropping — were applied to the training subset.

3.4.Dataset Preparation

Once preprocessed, the dataset was organized into its four grade-wise categories and split into training, validation, and testing subsets. The training subset was used to update the network parameters, the validation subset was used to track training progress and to guard against overfitting, and the testing subset was withheld entirely from training and used solely for the final performance evaluation reported in Section IV.

3.5.Deep Learning Model Development

Three convolutional neural network architectures — AlexNet, GoogLeNet, and ResNet50 — were adapted for this task using transfer learning. Each network was initialized with weights learned from the ImageNet dataset, which allowed the models to start from a set of generic, already-useful visual features rather than learning from scratch. The final classification layer of each network was replaced to output four class probabilities, one for each arecanut grade, and all three models were optimized using the Adam optimizer together with the categorical cross-entropy loss function. During fine-tuning, each network adjusted its learned representations to capture the texture, colour distribution, and surface characteristics that separate the four grades.

3.6.AlexNet

AlexNet was included in the comparison for its comparatively shallow but well-established architecture, consisting of five convolutional layers, three fully connected layers, and a final Softmax classifier. ReLU activations allow the network to model non-linear relationships between texture and quality, while dropout regularization is used to reduce overfitting, and max-pooling layers progressively reduce the spatial dimensions of the feature maps. The network was fine-tuned with the Adam optimizer at a learning rate of 0.0001. The class probability produced by AlexNet is expressed as:

$$\mathbf{P(yi) = Softmax\ (W3 \cdot ReLU\ (W2 \cdot ReLU(W1 * X' + b1) + b2) + b3)}$$

where the normalized image $X' = X / 255$ is passed through the learned convolutional filters W1, W2, and W3 with ReLU activation at each stage, Softmax converts the resulting logits into class probabilities, and categorical cross-entropy is used to measure the discrepancy between the predicted and true labels during training [16-20].

3.7.GoogLeNet

GoogLeNet (Inception V1) was chosen for its use of Inception modules, which apply parallel convolutional filters of different sizes — 1×1, 3×3, and 5×5 — together with max-pooling and Global Average Pooling (GAP) within a single module. This design allows the network to capture multi-scale visual features relevant to arecanut grading while keeping the number of trainable parameters relatively modest compared with an equivalently deep conventional CNN. Combined with transfer learning from ImageNet and the same augmentation strategy used for the other models, GoogLeNet produces its final prediction as:

$$\mathbf{Output = Softmax\ (W \cdot GAP\ (\Sigma\ InceptionModule\ (X)) + b)}$$

where the input X is processed through a sequence of Inception modules whose outputs are summed and reduced by Global Average Pooling; the pooled features are then passed through a dense layer with weights W and bias b before Softmax converts the resulting logits into class probabilities.

3.8.ResNet50

ResNet50 was selected as the deepest architecture in the comparison because its residual connections help maintain stable gradient flow even as network depth increases, which can otherwise make very deep networks difficult to train. In this implementation, features extracted by the ResNet50 backbone are reduced using Global Average Pooling and passed through a dense layer of 512 neurons with ReLU activation before reaching a final Softmax layer. Dropout and the same data-augmentation strategy used for the other two networks help control overfitting, and the network was again optimized with Adam at a learning rate of 0.0001. The resulting class probability is given by:

$$\mathbf{P(y) = Softmax\ (W2 \cdot Dropout\ (ReLU\ (W1 \cdot GAP(ResNet50(X)) + b1)) + b2)}$$

where the input image X is processed by the ResNet50 backbone for deep feature extraction, GAP reduces the resulting feature maps to a fixed-length vector, W1 and b1 define the dense ReLU layer, dropout with a rate of 0.5 helps prevent overfitting, and W2 and b2 in the final Softmax layer convert the output into class probabilities.

3.9.Performance Evaluation

The trained models were evaluated using four standard classification metrics — accuracy, precision, recall, and F1-score — supplemented by class-wise confusion matrices that reveal which grade pairs are most often confused with one another. Accuracy captures the overall fraction of correctly classified samples across all four grades:

$$\mathbf{Accuracy = (TP + TN) / (TP + TN + FP + FN)}$$

Precision measures how many of the samples predicted as belonging to a given grade were correctly assigned to it:

$$\mathbf{Precision = TP / (TP + FP)}$$

Recall measures how many of the samples that actually belong to a given grade were correctly identified:

$$\mathbf{Recall = TP / (TP + FN)}$$

and F1-score combines the two into a single balanced measure that is particularly useful when class sizes are unequal:

$$\mathbf{F1\text{-}Score = 2 \times (Precision \times Recall) / (Precision + Recall)}$$

In addition to these aggregate metrics, class-wise confusion matrices were generated for each of the three trained models so that grade pairs prone to misclassification could be identified and used to guide future dataset refinement.

3.10. Streamlit Deployment

To move the trained model out of the research environment and into a form that a processing unit could actually use, the best-performing classifier was integrated into a web application built with the Streamlit framework. The application allows a user to upload an arecanut image directly through the browser and receive an instant grade prediction, accompanied by a confidence score that indicates how certain the model is about that prediction. This real-time, browser-based interface was designed specifically to make the system usable by non-technical staff on the processing floor.

3.11. Explainable AI using Grad-CAM

Because a CNN's internal decision process is not directly visible to a human user, Gradient-weighted Class Activation Mapping (Grad-CAM) was incorporated into the deployed system to make each prediction more transparent. For a given input image, Grad-CAM produces a heatmap that highlights the spatial regions that contributed most strongly to the predicted grade, with warmer colours marking areas of higher importance and cooler colours marking areas that had little influence on the outcome. By overlaying this heatmap on the original image, the system gives the end user a visual explanation of why a particular grade was assigned, which helps build trust in the model's output and makes it easier to spot cases where the network may be relying on unexpected or irrelevant image regions.

4. Results and Discussion

4.1.Model Performance Comparison

Each of the three CNN architectures was evaluated on the held-out test set using accuracy, precision, recall, and F1-score, and the resulting values are summarized in Table 1. AlexNet delivered the strongest result across every metric, reaching an overall test accuracy of 99.15 percent. Out of 355 validation samples, only three were misclassified, which indicates that the network was able to resolve the visual differences between grades with a high degree of reliability.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4676-4683
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0614

Table 1 Comparative Model Performance

Model	Accuracy	Precision	Recall	F1-Score
AlexNet	99.15 %	99.17%	99.15 %	99.15 %
GoogLeNet	98.31 %	98.30%	98.31 %	98.29 %
ResNet50	98.31 %	98.35%	98.31 %	98.31 %

GoogLeNet and ResNet50 performed almost identically to one another, both settling at 98.31 percent accuracy — competitive figures, but slightly below AlexNet's result. This outcome is most plausibly explained by the relationship between dataset size and model capacity: with 1,792 images spread across four classes, the dataset may simply not be large enough to let the additional representational capacity of GoogLeNet or ResNet50 translate into a measurable accuracy gain, whereas AlexNet's comparatively compact parameter space appears to generalize slightly better under these data constraints.

Table 2 Performance Metrics of the Deep Learning Models (Normalized)

Metric	AlexNet	GoogLeNet	ResNet50
Accuracy	0.9915	0.9831	0.9831
Precision	0.9922	0.9817	0.9857
Recall	0.9915	0.9851	0.9852
F1-Score	0.9918	0.9830	0.9853

Table 2 restates these figures as normalized decimal values and confirms the same ordering seen in Table 1: AlexNet leads narrowly on every metric, with GoogLeNet and ResNet50 trailing by margins that are small enough to be considered practically comparable to one another.

4.2. Confusion Matrix Analysis

Class-wise confusion matrices were generated for AlexNet, GoogLeNet, and ResNet50 to examine classification behaviour at the level of individual grades. Across all three architectures, the overwhelming majority of test samples were assigned to their correct grade category, and the small number of errors that did occur were not spread evenly across all grade pairs. Instead, misclassifications were concentrated almost entirely between Rashidi and Gorabalu, the two grades whose colour and texture characteristics overlap the most. This pattern, consistent across all three models, points to a specific and addressable weakness in the current dataset rather than a general shortcoming of the classification approach, and it suggests that future work should prioritize collecting additional, more discriminative samples for precisely this grade pair.

4.3. Grad-CAM Visualization Analysis

To assess the quality of the explanations produced by the deployed system, Grad-CAM was applied to a sample boiled arecanut image that the AlexNet model classified as Gorabalu with a confidence score of 97.02 percent. The resulting heatmap showed the model's attention concentrated on the arecanut's surface texture, colour distribution, and structural boundaries — precisely the visual cues a human grader would use — while assigning little importance to the plain background surrounding the nut. This alignment between the model's learned attention and the visual cues that are actually relevant to grading offers reassurance that the network's high accuracy is not the product of some unrelated artefact in the images, but reflects genuine, grade-relevant feature learning.

4.4. Streamlit-Based Deployment Discussion

The Streamlit web application built around the trained AlexNet model performed as intended during testing, returning a predicted grade and confidence score immediately after an image was uploaded. Because the interface requires no specialized software or training to operate, it is realistic to imagine it being used directly by processing-line staff rather than remaining a research prototype. Paired with the Grad-CAM overlay, the deployed system gives users both a fast decision and a visual justification for that decision, which is a combination rarely offered by conventional manual grading.

Conclusion

This paper presented an automated grading system for boiled arecanut built on transfer-learning-based deep learning. A dedicated dataset of 1,792 images spanning four commercially recognized grades — Hasa, Bette, Rashidi, and Gorabalu — was collected and used to train and evaluate three CNN architectures. Among AlexNet, GoogLeNet, and ResNet50, AlexNet emerged as the strongest performer, reaching a test accuracy of 99.15 percent, and its predictions were made accessible through a Streamlit web application and interpretable through Grad-CAM visualization. Together, these components significantly reduce the dependence on manual grading, improve consistency across processing batches, and offer a practical, scalable path toward intelligent quality assessment in arecanut processing environments.

Future Enhancement

Future work will focus on enlarging the training dataset to capture a wider range of visual variation within each grade, particularly for the Rashidi and Gorabalu categories where the current models show the most confusion. A mobile-based version of the grading application would extend the system's usability to field-level deployment, where internet access or dedicated computing hardware may be limited. Finally, exploring more recent architectures such as EfficientNet and Vision Transformers may push classification accuracy and computational efficiency further beyond what has been demonstrated here.

Acknowledgment

The authors gratefully acknowledge Shivamogga district, Karnataka, for providing the boiled arecanut image dataset and for permitting controlled image acquisition under standardized conditions, without which this work would not have been possible. The authors also thank the faculty members and academic advisors of the Department of MCA, PES College of Engineering, Mandya, for their technical guidance and mentorship throughout this research. Appreciation is further extended to the open-source communities behind PyTorch, Streamlit, OpenCV, and Matplotlib, whose freely available tools made the implementation and deployment of the proposed grading system possible, and to the institution for providing the computational resources and facilities needed to carry out this work.

References

[1]. Ngamklet Nonglamin and Khwairakpam Amitab, "Grading of Arecanuts using Machine Learning Techniques," IEEE INSPECT Conference, 2025.
[2]. Avinash Kumar K. M., et al., "Transfer-Learning Based Multi-Class Areca Nut Image Classification under Uncontrolled Lighting on a Conveyor System," IEEE ICEI, 2024.
[3]. Adesh N. D., et al., "Segregation of Dehusked Arecanut using Artificial Intelligence Technique on Raspberry Pi," IEEE NKCon, 2023.
[4]. Chinmai Shetty, et al., "Automated Sustainable Grading of Chali Arecanuts: A Machine Learning Approach to Quality Assurance," IEEE DISCOVER, 2025.
[5]. Dhanush Ghate, et al., "Enhancing Arecanut Farming Profits through Technological Advancements: A CNN-Based Approach for Efficient Grading and Sorting," ICMCSI, 2024.
[6]. Ajit Danti and Suresha M., "Arecanut Grading Based on Three Sigma Controls and SVM," ICAESM, 2012.
[7]. Ajit Salunke and Sunilkumar Honnungar, "Estimation of Physical Properties of Unboiled Type Arecanuts using Load Cell and Ellipsoid Approximation," ICMCR, 2025.
[8]. Siddesha S., S. K. Niranjan, and V. N. Manjunath Aradhya, "Texture Based Classification of Arecanut," IEEE Conference, 2015.
[9]. Siddesha S., S. K. Niranjan, and V. N. Manjunath Aradhya, "Color Features and KNN in Classification of Raw Arecanut Images," IEEE Conference, 2018.
[10]. Bharadwaj N. K. and Dinesh R., "Possible Approaches to Arecanut Sorting/Grading using Computer Vision: A Brief Review," ICCCA, 2017.
[11]. Lokkhesh Annamalai and P. Vinothiyalakshmi, "Image-Based Fruit Quality Assessment using Deep Learning Algorithms," ICACRS, 2024.
[12]. Nidhi Kundu, Sonam, Geeta Rani, and Vijaypal Singh Dhaka, "A Deep Learning Based System for Automatic Sorting and Quality Grading of Citrus Fruits," ETNCC, 2024.
[13]. Yamani Chandana, et al., "A Lightweight

Deep Learning Framework for Real-Time Tomato Quality Grading Using Transfer Learning," IEEE GCAT, 2025.
[14]. Shih-Lun Wu, Hsiao-Yen Tung, and Yu-Lun Hsu, "Deep Learning for Automatic Quality Grading of Mangoes: Methods and Insights," IEEE ICMLA, 2020.
[15]. Shreedhara N. Hegde and H. S. Vijaya Kumar, "Karnataka Malenadu Region Arecanut Leafspot Disease Severity Assessment Using Machine Learning," International Conference on Smart Systems for Applications in Electrical Sciences (ICSSES), 2025.
[16]. Sneha Chikkanayakanahalli Siddalingadevaru, et al., "Disease Diagnosis and Sustainable Prevention in Arecanut Farming," IEEE ICEARS Conference, 2025.
[17]. Madhu B. G., Ram Kumar G., and Shreehari S. Rao, "Detection of Diseases in Arecanut Using Convolutional Neural Network," IEEE ICAIT Conference, 2024.
[18]. Pallavi Shetty, et al., "ArecaShield: AI Powered Arecanut Disease Detection and Advisory System," IEEE ICECST Conference, 2025.
[19]. Abhiram K., et al., "Motorized Arecanut Harvester using RF-based Control Interface," IEEE ICIMIA Conference, 2025.
[20]. Shih-Lun Wu, Hsiao-Yen Tung, and Yu-Lun Hsu, "Deep Learning for Automatic Quality Grading of Mangoes: Methods and Insights," IEEE ICMLA, 2020.