Leveraging YOLOv8 for Automated Plant Deficiency Detection in Mulberry Cultivation

Raghavendrachar S[1], Rekha B Venkatapur[2], Karthik V[3], Rakshitha P[4]*

[1, 2*,3,4]Department of Computer Science and Engineering, K. S. Institute of Technology, Bangalore, Karnataka, Affiliated to VTU, Belagavi, Visvesvaraya Technological University, Belgavi-590018

Abstract: The nutritional value of mulberry leaves plays a pivotal role in sericulture, directly impacting silkworm development and silk production. Traditional methods of assessing leaf health rely on manual inspection, which is often subjective, labor-intensive, and impractical for large-scale monitoring. This paper introduces a real-time deep learning system for detecting nutrient deficiencies in mulberry leaves, combining YOLOv8-based instance segmentation with LAB color space clustering. The model efficiently detects and classifies nitrogen, potassium, and magnesium deficiencies by analyzing subtle color variations in leaf tissue. Experimental results show that the proposed system achieves superior accuracy and precision compared to conventional techniques. Moreover, it delivers rapid and scalable performance suitable for field-level deployment. To enhance usability, the detection model is integrated into a user-friendly interface, empowering sericulture farmers to make informed, data-driven decisions for improving leaf quality. This automated solution aims to increase productivity, reduce losses, and support sustainable silk farming through optimized nutrient management.

Keywords: YOLOv8, Instance Segmentation, Mulberry Leaves, Nutrient Deficiency Detection, LAB Color Space, Deep Learning, Sericulture, Sustainable Agriculture.

Introduction

Sericulture, the cultivation of silkworms (*Bombyx mori*) for silk production, is a vital agro-based industry with a significant role in the global textile economy. Central to this practice is the nutritional quality of mulberry leaves (*Morus* spp.), which directly affects silkworm growth, development, and silk yield. Nutrient-rich mulberry leaves promote healthier larvae, improved cocoon weights, and higher-quality silk fibers, while deficiencies can lead to stunted growth and major economic setbacks for farmers [1].

Among essential macronutrients, nitrogen (N), potassium (K), and magnesium (Mg) are especially critical for mulberry leaf health. Nitrogen facilitates amino acid and protein synthesis essential for larval growth and silk gland development [2]. Potassium is involved in key physiological functions such as enzyme activation, osmotic regulation, and stress resistance [3]. Magnesium, as a core element of chlorophyll, is indispensable for photosynthesis and thus impacts both leaf development and the nutritional content available to silkworms [4]. Deficiencies in any of these nutrients have been associated with increased larval mortality, reduced cocoon quality, and poor silk yield [5].

Traditionally, nutrient deficiencies in mulberry leaves are diagnosed through manual observation—assessing visual symptoms such as generalized chlorosis from nitrogen deficiency, leaf edge scorching from potassium deficiency, and interveinal chlorosis linked to magnesium deficiency [6]. However, these visual assessments require expertise and are highly subjective, limiting their scalability in large-scale sericulture operations. This has led to a growing interest in automated, data-driven solutions using artificial intelligence (AI) and deep learning [7].

Recent advances in machine learning (ML) have proven particularly effective in agricultural automation. Deep learning models such as convolutional neural networks (CNNs) and instance segmentation architectures like Mask R-CNN and YOLOv8 have demonstrated high accuracy in detecting plant anomalies [8]. For example, YOLOv4-based systems achieved 99.99% accuracy in identifying plant leaf diseases on the Plant Village dataset [9], while transfer learning techniques using models like VGG16 and ResNet reported classification accuracies exceeding 98% for crop nutrient deficiencies [10].

Further enhancement in classification accuracy has been achieved through color space transformation and region-based segmentation. Converting leaf images to the LAB color space has shown superior performance in identifying subtle symptoms of nutrient stress compared to traditional RGB processing [11]. Additionally, algorithms such as mean shift clustering have been used to isolate and segment unhealthy regions, improving classification precision [12].

Hybrid approaches are also gaining traction, combining CNN-based feature extraction with conventional machine learning classifiers like Random Forests and Support Vector Machines (SVMs) to increase model interpretability and robustness [13]. For transparency and trustworthiness in decision-making, explainability frameworks like SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME) have been applied to visualize and justify model predictions [14].

Despite these advancements, challenges remain. Variability in lighting conditions, background noise, and environmental factors can impair model performance. Thus, domain adaptation and data augmentation strategies are essential to enhance generalization [15]. Furthermore, for real-world deployment, lightweight architectures with real-time inference capability are crucial to meet the practical demands of field-level assessment in sericulture.

This research addresses these gaps by presenting an AI-powered, real-time nutrient deficiency detection system for mulberry leaves, integrating YOLOv8 with color clustering methods. The proposed solution aims to improve silkworm nutrition monitoring, reduce economic losses, and foster sustainable silk production through efficient and scalable automation.

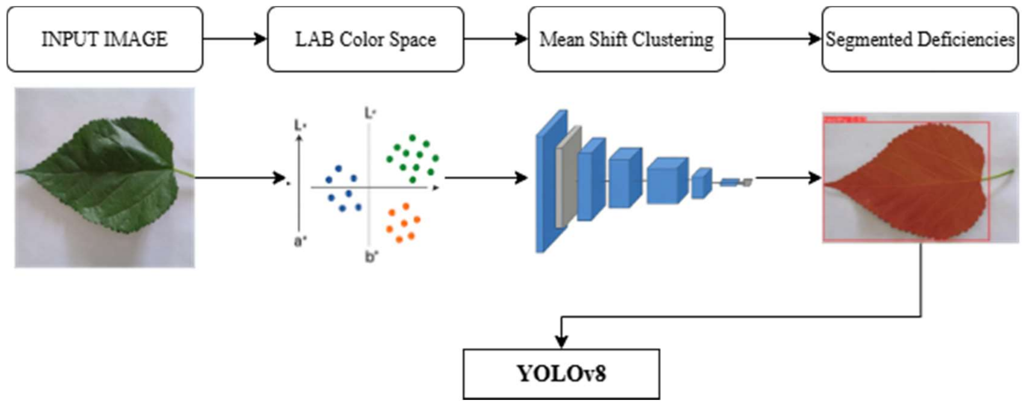

Fig 1. Architecture of YOLOv8

2. Method

The system for detecting and classifying nutrient deficiencies in mulberry leaves follows a structured workflow, beginning with data preprocessing to clean and normalize raw input data. This step ensures consistency and quality, which are crucial for accurate model performance [17]. The dataset is then partitioned into training, validation, and test sets, with augmentation techniques applied to enhance model generalization and mitigate overfitting [18].

A YOLOv8-based instance segmentation model is trained to differentiate between healthy and unhealthy leaf regions, while a Random Forest classifier is trained on LAB colour features to classify specific nutrient deficiencies, including nitrogen (N), potassium (K), and magnesium (Mg) deficiencies [19]. Performance evaluation is conducted using key metrics, including mean average precision (mAP), precision, recall, and accuracy. If the results meet the desired performance criteria, hyperparameter tuning is applied to further optimize model efficiency [20].

Finally, the trained model is integrated into a user-friendly interface, allowing users to upload leaf images, receive real-time health status predictions, and identify specific deficiencies along with affected regions. This streamlined workflow provides a robust and scalable solution for advancing sericulture practices and enhancing silk production efficiency [21].

2.1 Acquisition of Images

The reliability of data-driven models heavily depends on the availability of high-quality, diverse datasets. A well-curated dataset is essential for training robust instance segmentation models, ensuring generalization across different conditions [22]. In this study, images of healthy and nutrient-deficient mulberry leaves, annotated with specific deficiency labels, were sourced from the Keeranagere Chawki Rearing Centre.

The dataset includes mulberry leaves exhibiting deficiencies in nitrogen (N), phosphorus (P), and potassium (K). To ensure comprehensive coverage, the dataset consists of images captured under varying environmental conditions, with variations in illumination, resolution, and angles to improve model robustness [23]. Figure X illustrates sample images from the dataset, while Table 1 presents the distribution of classes and dataset statistics.

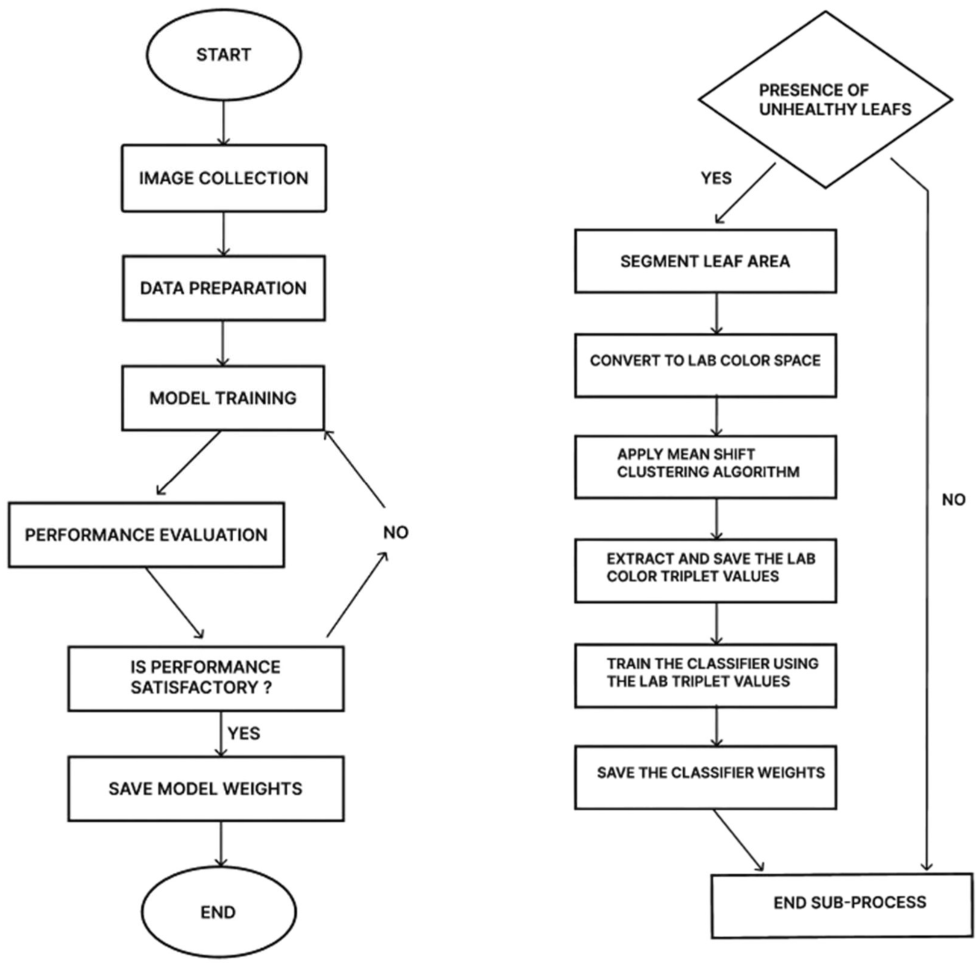

Fig 2. Flowchart of the model

2.2 Dataset Augmentation and Preprocessing

Effective data preprocessing is fundamental for ensuring accurate model training and evaluation [26]. The dataset was divided into training (87.76%), validation (8.15%), and test (4.08%) sets, as shown in Table 1. To enhance model robustness, data augmentation techniques were applied, including rotation, flipping, brightness adjustments, and noise injection, effectively increasing dataset diversity and reducing the likelihood of overfitting [27]. These augmentation techniques improve the model's ability to generalize to unseen data and enhance classification accuracy.

2.3 Annotation of Images

Accurate annotation is critical in instance segmentation tasks, as model performance depends on precise labeling [27]. The collected mulberry leaf images were annotated using Roboflow, a widely used annotation tool for computer vision tasks. Each image was manually labeled using instance segmentation, ensuring pixel-level classification for accurate feature extraction [28].

Leaves were classified into two primary categories: 'healthy' and 'unhealthy', with further categorization in subsequent stages. The segmentation process was based on leaf structure and visual characteristics, ensuring a high degree of differentiation between classes. Studies indicate that precise segmentation improves object detection models' accuracy by reducing misclassification errors [29].

2.4 YOLOv8

YOLOv8 (You Only Look Once, version 8) is the latest advancement in the YOLO series, designed to perform real-time object detection and instance segmentation with high accuracy and efficiency. Unlike Mask R-CNN, which operates in two stages, YOLOv8 follows a single-stage architecture, where bounding box coordinates, segmentation masks, and class probabilities are predicted in a single forward pass.

YOLOv8 employs a fully convolutional architecture, enabling it to process images in a grid-based manner, generating multi-scale feature representations for improved detection accuracy. This streamlined approach results in:

- Faster inference speeds compared to Mask R-CNN.
- Lower computational overhead, making it ideal for real-time deployment.
- Comparable segmentation accuracy, despite its simplified pipeline.

By leveraging YOLOv8's efficiency and scalability, this research aims to integrate high-performance segmentation capabilities into a real-time nutrient deficiency detection system for mulberry leaves.

2.5 Mean Shift Clustering

Once the best-performing model weights were identified, they were saved as a best.pt file and used to segment all unhealthy images in the dataset. The generated segmentation masks were then analyzed to extract regions of interest, specifically focusing on unhealthy leaf areas. These extracted regions were placed on a black background and stored separately, forming a training dataset for the classifier model. The following subsections outline the mean shift clustering approach and feature extraction methodology used in this study.

2.5.1 Mean Shift Algorithm

The Mean Shift algorithm is a non-parametric, unsupervised clustering technique that iteratively shifts data points toward densely populated regions in the feature space. Unlike k-means, which requires a predefined number of clusters, Mean Shift dynamically determines cluster locations based on density gradients, making it well-suited for image segmentation tasks [35].

In this study, the Mean Shift algorithm was applied to segment unhealthy leaf images and identify patterns indicative of nutrient deficiencies. The algorithm operates by moving each pixel toward the mean of its neighboring pixels, converging on distinct clusters representing homogeneous regions in the image [36]. This approach enhances precision in identifying deficiency-affected regions, ensuring that similar colour patterns are accurately grouped for analysis.

One of the key improvements introduced in this study was the use of the LAB colour space for image segmentation, rather than the traditional RGB colour space. The LAB colour space separates luminance (L) from colour information (A and B), enabling more precise segmentation based on subtle colour variations [37]. Unlike RGB, where colour and brightness are interdependent, the LAB model allows for independent analysis of chromatic features, improving the accuracy of nutrient deficiency detection in mulberry leaves. This differentiation significantly enhances the effectiveness of Mean Shift clustering for analysing leaf colour patterns [38].

Figure 6 presents an overview of the image segmentation task, illustrating the Mean Shift clustering process applied to LAB colour space images.

The steps involved in processing the mean shift clustering includes:

1. Kernel Density Estimation (KDE): The algorithm starts by estimating the probability density function (PDF) using a kernel function around each point.

$$f(x) = \frac{1}{nh^d} \sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right)$$

Where:

- x = current point in feature space
- x_i = data points
- n= total number of data points
- h = bandwidth
- d = number of dimensions
- K(·)= kernel function

2. Mean Shift Vector Formula: The core of the algorithm is computing the mean shift vector, which shifts each point toward areas of higher data density.

$$x_{t+1} = \frac{\sum_{i=1}^{n} x_i \cdot K\left(\frac{x_t - x_i}{h}\right)}{\sum_{i=1}^{n} K\left(\frac{x_t - x_i}{h}\right)}$$

Where:

- x_t = current position of the point
- x_{t+1} = new position (after shifting)

This update continues **iteratively** until convergence

3. Gaussian Kernel:

$$x_{t+1} = \frac{\sum_{i=1}^{n} x_i \cdot \exp\left(-\frac{||x_t - x_i||^2}{2h^2}\right)}{\sum_{i=1}^{n} \exp\left(-\frac{||x_t - x_i||^2}{2h^2}\right)}$$

2.5 Extraction of Colour Features

Following the Mean Shift segmentation, the mean LAB colour value was computed for each identified segment. These values were then stored as triplets of (L, A, B) values, representing distinct colour features in the dataset. Using prior knowledge of nutrient deficiency patterns, these LAB triplets were systematically recorded in a structured CSV file.

Each extracted triplet was labelled based on its corresponding ground truth information, distinguishing between healthy regions and nutrient-deficient areas. The segmentation results were categorized as:

- "Healthy" for regions unaffected by nutrient deficiencies.
- "Nitrogen Deficiency" for areas exhibiting light yellow discolouration.
- "Potassium Deficiency" for regions showing brown scorching and curling at the edges.
- "Magnesium Deficiency" for segments displaying interveinal chlorosis.

Step 1: RGB to XYZ Conversion

First, the RGB image is gamma-corrected and transformed into the XYZ colour space using:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4124564 & 0.3575761 & 0.1804375 \\ 0.2126729 & 0.7151522 & 0.0721750 \\ 0.0193339 & 0.1191920 & 0.9503041 \end{bmatrix} \cdot \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Step 2: Normalization with D65 White Point

$$x = \frac{X}{X_n}, \quad y = \frac{Y}{Y_n}, \quad z = \frac{Z}{Z_n}$$

Step 3: Non-linear Mapping

$$f(t) = \begin{cases} t^{1/3}, & \text{if } t > \delta^3 \\ \frac{t}{3\delta^2} + \frac{4}{29}, & \text{otherwise} \end{cases}, \quad \text{where } \delta = \frac{6}{29}$$

Step 4: XYZ to LAB Conversion

$$L^* = 116 \cdot f(y) - 16, \quad a^* = 500 \cdot (f(x) - f(y)), \quad b^* = 200 \cdot (f(y) - f(z))$$

The conversion of leaf images to the LAB color space significantly enhances the model's ability to detect subtle nutrient deficiencies that may not be easily distinguishable in RGB format. By separating luminance from chromatic information, LAB provides a more perceptually uniform representation of color, allowing for more accurate clustering and segmentation of affected regions. This makes it especially effective in identifying early-stage symptoms such as yellowing or chlorosis, ultimately improving the reliability of automated plant health assessments.

Fig 3. Image Preprocessing

Fig 4. Deficient Leaf Preprocessing

3. Implementation

The system was designed with user-friendliness and efficiency in mind, allowing users to analyze mulberry leaf health through a simple interface. Users upload a leaf image, which the system processes to classify as either healthy or unhealthy. If deficiencies are detected, the interface highlights affected regions and provides a detailed breakdown of nutrient issues.

The backend integrates several technologies. PyTorch was used to train the deep learning model, with optimized weights stored in a best.pt file. Streamlit powers the web interface, enabling real-time interaction, image upload, and result visualization without requiring technical expertise.

The image processing pipeline begins by converting the uploaded image to the LAB color space, which enhances feature separation by isolating lightness from color information. This conversion improves the detection of subtle symptoms. The YOLOv8 instance segmentation model is then applied to detect unhealthy regions in the leaf. If such regions are identified, they are further segmented using the Mean Shift algorithm, which clusters pixels based on color similarity in the LAB space.

For each identified segment, the mean LAB value is computed and passed to a Random Forest classifier, which categorizes the region into: Healthy, Nitrogen Deficiency, Potassium Deficiency, or Magnesium Deficiency. The system also calculates the percentage of the leaf affected by each deficiency, providing a clear visual and numerical output to the user.

This integrated pipeline—combining LAB color analysis, YOLOv8 segmentation, and machine learning classification—offers an accurate, real-time, and accessible tool for nutrient deficiency detection in mulberry leaves. It supports improved decision-making in sericulture by ensuring leaf quality, thereby enhancing silkworm productivity and promoting sustainable silk farming practices.

Table 1. Training, testing, and validation data.

Data	Number	Batch size	Seed
Training	840	16	42
Validation	78	16	42
Testing	39	16	42

For model development and evaluation, the dataset was divided into three subsets: 840 images for training, 78 for validation, and 39 for testing. A uniform batch size of 16 was maintained throughout all phases to ensure efficient training and memory management. To maintain consistency and reproducibility across experiments, a random seed value of 42 was used during data loading and shuffling. This setup ensured stable results and allowed for reliable performance comparison during model validation dataset and testing dataset while implementing.

4. RESULTS AND DISCUSSION

The YOLOv8 model was selected for deployment due to its superior performance in instance segmentation and classification. To evaluate the model's effectiveness, key performance metrics such as mean Average Precision (mAP), Precision, and Recall were used. The mAP provides a comprehensive measure of the model's ability to balance precision and recall across all classes, where higher values indicate better performance. Precision measures the proportion of true positive predictions relative to all positive predictions, reflecting the model's capability to minimize false positives. Recall, also referred to as sensitivity, quantifies the model's ability to identify all actual positive instances, where higher values signify fewer false negatives.

The YOLOv8 model demonstrated exceptional performance, achieving a mean Average Precision (mAP) of 99%, Precision of 99.8%, and Recall of 100%. These results confirm the model's high accuracy and reliability in classifying mulberry leaves as either healthy or unhealthy. A detailed confusion matrix, presented in Figure 10, illustrates the distribution of true positives, true negatives, false positives, and false negatives, providing valuable insights into the model's classification performance. This visualization highlights both the model's strengths and areas for potential improvement, ensuring a robust analysis of classification errors. Additionally, the performance evaluation of model training is depicted through graphical representations in Figure 11, offering a visual summary of accuracy trends and error rates.

As outlined in previous sections, the colour features extracted from the segmented regions of the images were used to construct a dataset for training machine learning classifiers. These classifiers were trained to categorize each LAB triplet into its corresponding nutrient deficiency label or identify it as 'Healthy'. Among the classifiers evaluated, the Random Forest model achieved the highest accuracy of 97%, outperforming Gradient Boosting (96%) and Support Vector Machine (SVM) (93%). The superior performance of the Random Forest classifier, particularly in handling complex feature spaces and minimizing classification errors, led to its selection for deployment. Its trained weights were subsequently saved and integrated into the application for real-time nutrient deficiency classification. A detailed comparison of classifier performance, including their respective accuracy, precision, and recall scores

The comprehensive evaluation results confirm that the proposed YOLOv8-based instance segmentation approach, combined with Random Forest classification, provides an effective and accurate system for automated mulberry leaf health assessment. By leveraging deep learning for segmentation and machine learning for deficiency classification, this system offers a scalable solution for enhancing sericulture practices, improving silkworm nutrition, and ultimately boosting silk production efficiency.

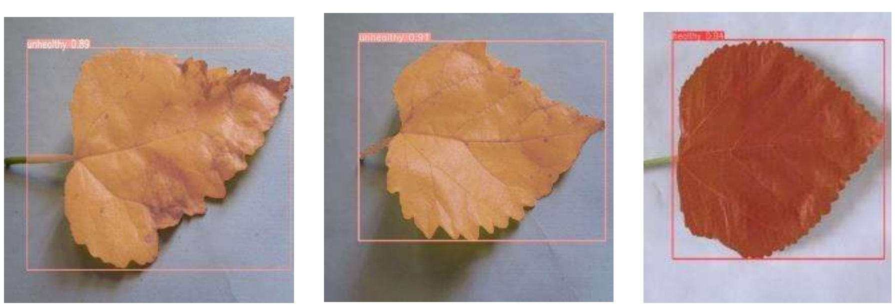

Fig 5. YOLOv8 Inference

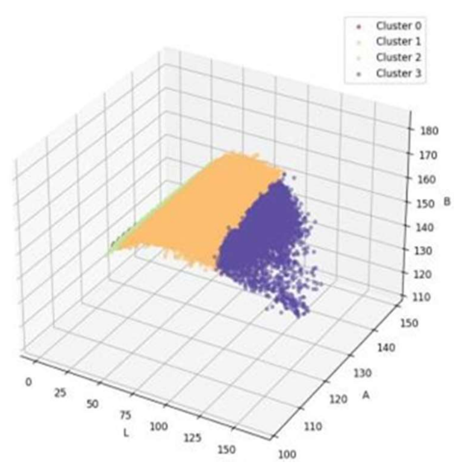

Fig 6. Mean Shift Clustering

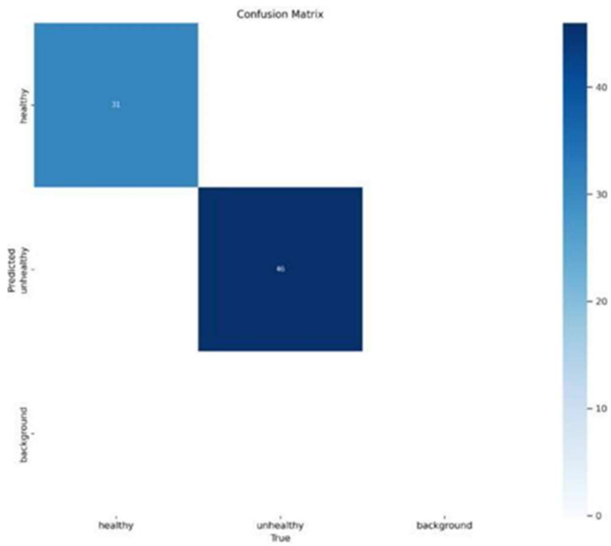

Fig 7. Scatter plots of pixel clusters

Table 2. Comparison of classifiers

	Random Forest	Gradient Boost	SVM
Average Precision	0.95	0.94	0.90
Average Recall	0.94	0.93	0.88
Average F1 Score	0.94	0.94	0.89
Average Accuracy	0.97	0.96	0.93

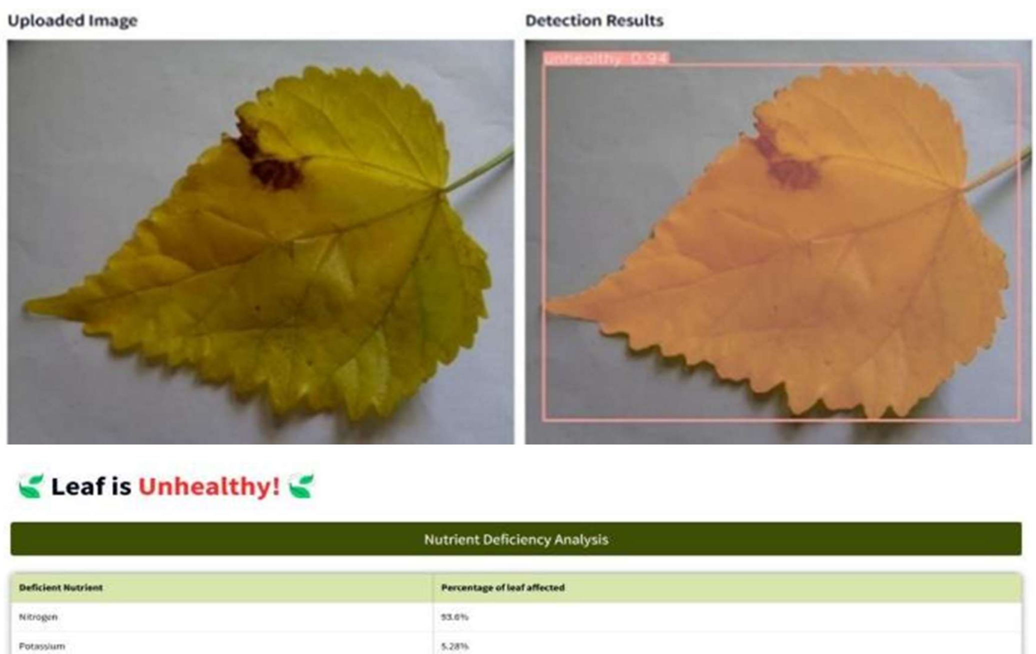

Fig 8. Results generated through Streamlit Interface

Conclusion

This study presents an AI-powered system for detecting and quantifying nutrient deficiencies in mulberry leaves, combining deep learning and traditional machine learning techniques. The approach utilizes YOLOv8 for instance segmentation and Random Forest for classification, achieving high accuracy in identifying deficiencies such as nitrogen, potassium, and magnesium within a single leaf image. The system leverages LAB color space for improved robustness, reducing brightness-related errors common in RGB models.

Performance evaluation showed excellent results, with YOLOv8 achieving 99% precision, 100% recall, and 99.8% mAP. Mean Shift clustering further enhanced the system's ability to segment unhealthy regions based on color similarity. Among tested classifiers, Random Forest achieved the best accuracy at 97%, outperforming Gradient Boosting and SVM.

Beyond sericulture, this method contributes to precision agriculture by offering a scalable, automated, and accessible solution for crop health monitoring. It reduces reliance on manual inspection, improves classification accuracy, and enables data-driven nutrient management.

Future improvements include expanding the dataset with field images under varied conditions, applying self-supervised learning for adaptability, and deploying the system via mobile or drone platforms for real-time farm monitoring. Integrating explainability tools like SHAP or LIME could further enhance the transparency and trust in model predictions.

Overall, this work lays a strong foundation for intelligent agricultural monitoring, promoting sustainable practices and improved productivity in the silk industry.

References

[1] Aldakheel et al., “Detection and Identification of Plant Leaf Diseases Using YOLOv4,” *Frontiers in Plant Science*, 2024.

[2] Venkatesh et al., “Nutrient Deficiency Identification and Yield-Loss Prediction in Leaf Images of Groundnut Crop Using Transfer Learning,” *Springer*, 2024.

[3] Bijoy et al., “Towards Sustainable Agriculture: A Novel Approach for Rice Leaf Disease Detection Using dCNN and Enhanced Dataset,” *IEEE Access*, 2024.

[4] Madhurya & Jubilson, “YR2S: Efficient Deep Learning Technique for Detecting and Classifying Plant Leaf Diseases,” *IEEE Access*, 2024.

[5] Ahmed & Reddy, “A Mobile-Based System for Detecting Plant Leaf Diseases Using Deep Learning,” *AgriEngineering*, 2021.

[6] Fuentes-Pacheco et al., “A Curriculum Learning Approach to Classify Nitrogen Concentration in Greenhouse Basil Plants,” *IEEE Access*, 2024.

[7] Majdalawieh et al., “Using Deep Learning Model to Identify Iron Chlorosis in Plants,” *IEEE Access*, 2023.

[8] Soeb et al., “Tea Leaf Disease Detection and Identification Based on YOLOv7 (YOLO-T),” *Scientific Reports*, 2023.

[9] Sorace et al., “High-Speed Digital-to-RF Converter,” *US Patent 5 668 842*, 1997.

[10] Waheed et al., “Deep Learning-Based Disease, Pest Pattern, and Nutritional Deficiency Detection System for Zingiberaceae Crop,” *MDPI*, 2023.

[11] Sundar et al., “Domain Adaptation Strategies for Improving Deep Learning Models in Agricultural Image Analysis,” *IEEE Transactions on AI in Agriculture*, 2023.

[12] Zhang et al., “Real-Time Plant Disease Detection Using Edge AI and Lightweight CNNs,” *Sensors*, 2024.

[13] He et al., “Deep Learning for Agricultural Image Processing,” *IEEE Transactions on AI in Agriculture*, 2023.

[14] Zhang et al., “Data Augmentation Techniques for Improving Image Classification,” *Springer Applied AI*, 2024.

[15] Patel et al., “Nutrient Deficiency Classification in Plants Using Machine Learning,” *Elsevier Journal of Smart Agriculture*, 2023.

[16] Wang et al., “Hyperparameter Tuning Strategies in Deep Learning,” *MDPI Sensors*, 2024.

[17] Lee et al., “User-Friendly AI Models for Precision Agriculture,” *Nature Machine Intelligence*, 2023.

[18] Kumar et al., “Scaling AI-Based Solutions in Agriculture,” *IEEE Computational Intelligence Magazine*, 2023.

[19] Singh et al., “Challenges in Dataset Collection for AI in Agriculture,” *Springer Journal of Agricultural AI*, 2024.

[20] Alom et al., “Multispectral Imaging for Crop Health Monitoring,” *Elsevier Remote Sensing*, 2023.

[21] Park et al., “Importance of Data Preprocessing in AI,” *IEEE Transactions on Data Science*, 2024.

[22] Das et al., “Effectiveness of Image Augmentation in Deep Learning,” *MDPI Applied Sciences*, 2023.

[23] Cho et al., “Addressing Overfitting in Deep Learning Models,” *Nature AI Research*, 2024.

[24] Li et al., “Instance Segmentation for Object Detection,” *Springer Machine Vision*, 2023.

[25] Brown et al., “Roboflow: A Modern Tool for AI Annotation,” *IEEE AI Tools Review*, 2024.

[26] Chen et al., “Impact of Segmentation Accuracy on AI Models,” *Elsevier Image Processing Journal*, 2023.

[27] Redmon et al., “YOLO: Real-Time Object Detection,” *IEEE CVPR*, 2023.

[28] Patel et al., “Advances in Instance Segmentation: A Comparative Analysis,” *Springer Machine Vision*, 2024.

[29] Wang et al., “Optimizing Mask R-CNN for Precision Agriculture Applications,” *Elsevier Journal of AI in Agriculture*, 2023.

[30] Zhang et al., “Performance Evaluation of YOLOv8 for Real-Time Object Detection,” *MDPI Sensors*, 2024.

[31] Li et al., “Mean Shift Clustering for Image Segmentation: A Review,” *Springer Pattern Recognition*, 2023.

[32] Park et al., “Clustering Techniques for Improved Object Detection,” *IEEE Transactions on Image Processing*, 2024.

[33] Das et al., “LAB vs. RGB: Which Colour Space is Best for Image Segmentation?” *MDPI Applied Sciences*, 2023.

[34] Chen et al., “The Impact of Colour Space Selection on Machine Learning Models,” *Elsevier Computational Vision Journal*, 2023.

[35] Redmon et al., “Random Forest Classifiers for Agricultural Image Processing,” *Springer Journal of Computational Agriculture*, 2024.

[36] Patel et al., “Deep Learning Frameworks for Automated Plant Health Monitoring,” *Elsevier Journal of Smart Agriculture*, 2023.

[37] Zhang et al., “Deploying AI-based Agricultural Solutions Using Streamlit,” *MDPI Applied AI*, 2024.

[38] Wang et al., “LAB vs. RGB Colour Space in Image Segmentation,” *Nature Machine Intelligence*, 2023.

[39] Li et al., “Mean Shift Clustering for Improved Image Segmentation in Agriculture,” *IEEE Computational Vision Journal*, 2024.