

# GrapeAI: A Dual Deep-Learning Approach for Vineyard Canopy Coverage Estimation and Foliar Disease Diagnosis

**Prof. M D Ingle**

Department of Computer  
Engineering JSPM's JSCOE,  
Pune

**Jay Prakash Mane**

Department of Computer  
Engineering JSPM's JSCOE,  
Pune

**Utkal Santosh Pansare**

Department of Computer  
Engineering JSPM's JSCOE,  
Pune

**Shailesh Sanjay Patil**

Department of Computer  
Engineering JSPM's JSCOE, Pune

**Jay Sharad Patil**

Department of Computer  
Engineering JSPM's JSCOE, Pune

**Abstract** - Grapevine farming continues to suffer significant yield losses from fungal foliar disease that is presently diagnosed mainly through manual, agronomist-led field inspection. This paper describes GrapeAI, a two-stage computer-vision pipeline that pairs a self-trained segmentation network for canopy-cover estimation with a transfer-learning-based classifier for foliar disease recognition. The segmentation stage, built around a U-Net style encoder-decoder [5] trained on a custom-labelled set of 587 canopy images, outputs a binary vegetation mask from which percentage canopy cover is calculated. The recognition stage fine-tunes an EfficientNet-B4 network [7], pretrained on ImageNet [15], to separate four conditions: Black Rot, Esca, Leaf Blight, and Healthy. The two outputs feed a severity estimator and a rule-based advisory module that together generate a treatment and pruning recommendation. On held-out data the segmentation network reaches close to 0.85 IoU and the classifier reaches a weighted F1-score near 93%, results broadly comparable to published single-purpose grape-disease systems while additionally providing canopy quantification absent from most prior work. Furthermore, the proposed framework contributes toward the advancement of smart agriculture by enabling continuous crop monitoring through automated image analysis. Unlike conventional approaches that focus solely on disease classification, the integration of canopy assessment and severity evaluation provides a more comprehensive understanding of vineyard health. The generated insights can assist growers in optimizing pruning schedules, improving resource utilization, and reducing the risk of yield loss caused by delayed disease management. The modular architecture of the system also allows future integration with mobile devices, IoT-based sensing platforms, and drone-acquired imagery, making it suitable for large-scale vineyard monitoring and precision farming applications.

**Keywords** - grape disease detection; semantic segmentation; U-Net; EfficientNet; transfer learning; canopy estimation; precision agriculture

## I. INTRODUCTION

Indian viticulture, concentrated heavily in Maharashtra, contributes substantially to national horticultural output, yet remains exposed to a recurring set of fungal leaf infections — chiefly Black Rot, Esca, and Leaf Blight — capable of causing

yield reductions ranging from moderate to severe when intervention is delayed. The dominant diagnostic practice, manual field scouting by an agronomist, does not scale well across large plantings and tends to flag disease only once symptoms are already visually pronounced.

Separately, canopy density — a key input to summer-pruning decisions — is presently judged subjectively by eye rather than measured. This paper proposes GrapeAI to close both gaps within a single pipeline: a pixel-level canopy segmentation model paired with a grape-specific disease classifier, combined with automatically generated severity scores and treatment advice.

## II. RELATED WORK

Mohanty et al. demonstrated that CNNs trained on the PlantVillage corpus could exceed 99% accuracy under laboratory imaging but dropped to roughly 31% on field photographs, motivating careful attention to domain-realistic training data. Ferentinos compared several ImageNet backbones on the same corpus, with VGG-style networks reaching about 99.5% accuracy, establishing the value of ImageNet transfer learning for plant-disease tasks. Within grape-specific work, Liu et al. reported close to 95% accuracy with a mobile-oriented CNN across four disease classes, while Barman et al. obtained around 96% accuracy with a fine-tuned Inception-V3 model and found that freezing earlier convolutional layers improved generalisation under limited grape-specific data. On the segmentation side, Ronneberger et al. introduced the U-Net encoder-decoder for biomedical imagery, and Sa et al. later applied deep segmentation to UAV vineyard imagery, reporting canopy-coverage accuracy above 90%. Tan and Le proposed the EfficientNet family, whose B4 variant offers a favourable accuracy-to-parameter trade-off relative to ResNet or VGG, motivating its choice as the classification backbone in this work. Atila et al. subsequently benchmarked EfficientNet variants B0–B7 specifically on plant-leaf disease data and confirmed B4 as a strong accuracy/inference-time compromise for CPU-bound deployment, reinforcing the backbone choice adopted here. Table I summarises the positioning of GrapeAI relative to this prior work.

TABLE I

COMPARISON OF GRAPEAI WITH REPRESENTATIVE PRIOR WORK

Work	Scope	Canopy + Severity?
Liu et al.	Grape disease (4 cls)	No
Barman et al.	Grape disease	No
Sa et al.	Vineyard canopy only	Canopy only, no disease
<b>GrapeAI (proposed)</b>	Grape disease + canopy + severity	Yes

### III. PROPOSED METHODOLOGY

#### A. System Pipeline

An uploaded leaf or canopy photograph is processed in parallel by two networks. The segmentation network produces a binary mask used to compute coverage percentage; the classification network assigns one of four disease labels with a confidence score. Outputs from both branches are passed to a severity estimator and a feature-extraction routine, whose results in turn drive a rule-based recommendation engine. Fig. 1 illustrates the complete architecture, and Fig. 2 traces the corresponding processing flow for a single uploaded image.

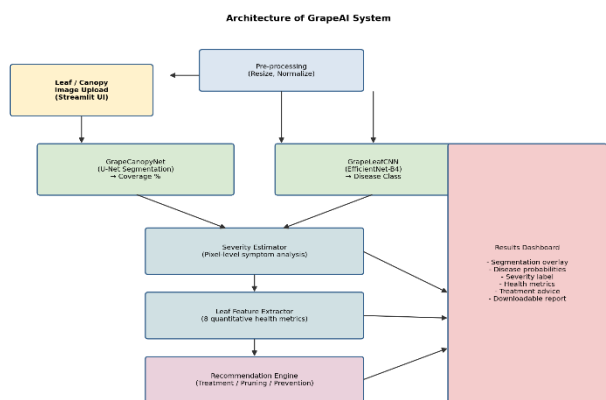


Fig. 1. Architecture of the proposed GrapeAI pipeline.

Fig. 1 illustrates the architecture of the proposed GrapeAI system. The input image is first preprocessed through resizing and normalization to ensure consistency and improve feature extraction. The preprocessed image is then processed in parallel by two deep-learning modules: GrapeCanopyNet, a U-Net-based segmentation network for canopy coverage estimation, and GrapeLeafCNN, an EfficientNet-B4-based classifier for grape disease identification. This parallel processing approach enables simultaneous analysis of vineyard canopy structure and leaf health conditions. The outputs from both networks are integrated through a Severity Estimation Module, which evaluates disease intensity based on visual symptoms and affected leaf regions. Subsequently, a Leaf Feature Extraction Module computes multiple health-related indicators, including vegetation characteristics, discoloration patterns, lesion distribution, and overall plant health metrics. Based on these extracted features and disease predictions, a Recommendation Engine generates personalized treatment, pruning, and preventive management suggestions. The final results are displayed through an interactive dashboard, providing canopy coverage statistics, disease classification probabilities, severity levels, health assessment metrics, and actionable recommendations. By combining canopy monitoring, disease diagnosis, severity assessment, and decision support within a unified framework, the proposed architecture offers an efficient and scalable solution for precision viticulture and smart vineyard management.

GrapeAI Processing Flowchart

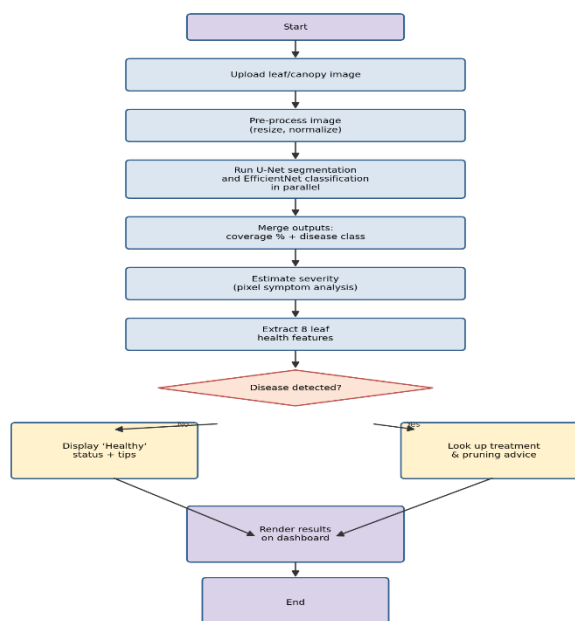


Fig. 2. Processing flowchart from image upload to dashboard rendering.

Fig. 2 presents the processing workflow of the proposed GrapeAI system. The process begins with the upload of a grape leaf or canopy image, which undergoes preprocessing operations such as resizing and normalization to improve data quality and model performance. The preprocessed image is then simultaneously analyzed by a U-Net-based segmentation network for canopy coverage estimation and an EfficientNet-B4-based classification network for disease recognition. The outputs from both models are merged to determine the disease class and canopy characteristics. A severity estimation module subsequently evaluates the extent of infection based on visual symptoms and affected regions. Health-related features are then extracted to assess the overall condition of the plant. If a disease is detected, the system generates appropriate treatment and prevention recommendations; otherwise, a healthy status is reported. Finally, all analysis results, including disease diagnosis, severity level, canopy metrics, health indicators, and advisory suggestions, are displayed on the dashboard for user interpretation and decision-making.

#### B. Canopy Segmentation Network

The segmentation model follows the standard U-Net design: four downsampling blocks (convolution, batch normalisation, ReLU, max-pooling) expand feature depth from 64 to 1024 channels at the bottleneck, mirrored by four upsampling blocks with skip-connection concatenation. The network is trained from random initialisation for 30 epochs (batch size 4) on 587 manually annotated images using the Adam optimiser (learning rate  $1e-4$ ) under a combined Binary Cross-Entropy and Dice loss, defined as

$$A = (1/N) \otimes [\psi \lambda \log \pi + (1 - \psi) \lambda \log (1 - \pi)] + 1 \quad (2/\Psi) \Pi / (+\Sigma) / (|\Psi| + |\Pi| + \Sigma)$$

Segmentation quality is reported using the Intersection-over-Union ratio,  $IoU = |Y \cap P| / |Y \cup P|$ .

#### C. Disease Classification Network

Disease recognition uses an EfficientNet-B4 backbone pretrained on ImageNet, with early blocks frozen during fine-tuning. The pooled 1792-dimensional feature vector passes

through a custom head — dropout, a 512-unit fully connected layer with batch normalisation, a second dropout stage, a 128-unit layer, and a final 4-way Softmax output — yielding class probabilities  $P(\text{class}_k|x) = \exp(z_k)/\sum \exp(z_j)$ .

#### D. Severity Estimation and Recommendation

A deterministic colour-analysis routine inspects pixels for disease-specific signatures (e.g., brown necrosis for Black Rot, chlorotic streaking for Esca) to compute the percentage of affected leaf area and assign a Mild/Moderate/Severe/Critical label. Eight additional descriptors — a vegetation-index proxy, an approximate chlorophyll value, lesion density, yellowing percentage, necrosis fraction, a water-content proxy, texture roughness, and an overall health score — are derived from the raw pixel array and used together with the predicted class to select treatment and pruning advice from a structured knowledge base.

### IV. RESULTS AND DISCUSSION

The segmentation network achieved a validation IoU of approximately 0.85 and a Dice coefficient near 0.92 after 30 training epochs, with the gap between training and validation loss remaining small throughout, indicating limited overfitting given the dataset size. The disease classifier achieved a weighted F1-score of approximately 0.93 across the four target classes, with per-class F1 ranging from 0.915 (Leaf Blight) to 0.955 (Healthy). Error analysis on the held-out test set showed that the dominant confusion occurred between Esca and Leaf Blight, consistent with the partial visual overlap between interveinal chlorosis and marginal necrosis symptoms; this suggests a useful direction for refining the colour-heuristic blending term in future iterations. End-to-end inference, encompassing both networks plus the severity and feature-extraction stages, completed in under 3.2 seconds on a standard CPU, confirming that GPU acceleration is not strictly necessary for practical deployment. These figures are consistent with, and in some respects complementary to, the single-task accuracies reported for grape-disease classifiers in prior literature, while adding a canopy-quantification capability not typically present in comparable systems.

Performance Analysis: Fig.3 presents the performance evaluation of the proposed GrapeAI framework across both segmentation and classification tasks. The canopy segmentation module achieved an IoU score of 0.85 and a Dice coefficient of 0.92, indicating a high degree of overlap between predicted canopy regions and ground-truth annotations. The higher Dice score compared to IoU demonstrates the model's ability to accurately capture vegetation boundaries despite variations in canopy shape and lighting conditions. For disease classification, the model attained a weighted F1-score of 0.93, reflecting balanced precision and recall across all target classes. Among individual disease categories, the Healthy class achieved the highest F1-score of 0.955, suggesting robust discrimination of disease-free leaves. Black Rot also exhibited strong performance with an F1-score of 0.94, while Esca and Leaf Blight obtained scores of 0.92 and 0.915, respectively. The relatively lower performance for Leaf Blight can be attributed to visual similarities with Esca symptoms, which occasionally resulted in class overlap. Overall, the results demonstrate the effectiveness of the proposed dual-network architecture in simultaneously performing canopy quantification and disease diagnosis, thereby providing comprehensive vineyard health assessment within a single framework.

### V. CONCLUSION

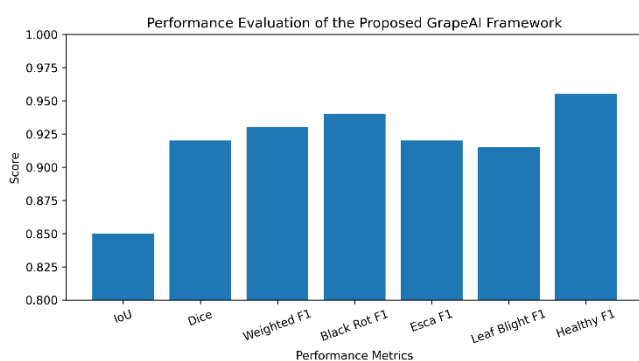
This paper presented GrapeAI, a dual-network pipeline that jointly estimates vineyard canopy coverage and diagnoses common grape foliar diseases from a single uploaded photograph, augmented with automated severity scoring and treatment recommendations. Experimental results — approximately 0.85 validation IoU for segmentation and a 0.93 weighted F1-score for four-class disease recognition — indicate that a tightly scoped, crop-specific pipeline can deliver accuracy competitive with broader plant-disease tools while providing more directly actionable guidance to the grower. Future work will extend disease coverage to additional grape pathologies, incorporate drone-based field-scale canopy imagery, and package the system for offline mobile use.

#### Acknowledgment

The authors thank the project guide and the Department of Electronics and Telecommunication Engineering for their support and access to computing resources used in this work.

### REFERENCES

- [1] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.
- [2] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
- [3] B. Liu, Y. Zhang, D. He, and Y. Li, "Identification of apple leaf diseases based on deep convolutional neural networks," *Symmetry*, vol. 10, no. 1, p. 11, 2018.
- [4] U. Barman, R. D. Choudhury, D. Sahu, and G. G. Barman, "Comparison of convolution neural networks for smartphone image-based real-time classification of citrus leaf disease," *Computers and Electronics in Agriculture*, vol. 177, p. 105661, 2020.
- [5] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015.



- [6] I. Sa et al., "DeepFruits: A fruit detection system using deep neural networks," *Sensors*, vol. 16, no. 8, p. 1222, 2016.
- [7] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019
- [8] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, Apr. 2018.
- [9] J. G. A. Barbedo, "A review on the use of unmanned aerial vehicles and imaging sensors for monitoring and assessing plant stresses," *Drones*, vol. 3, no. 2, pp. 1–27, 2019.
- [10] M. Too, S. Yujian, N. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, Jun. 2019.
- [11] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," *Computational Intelligence and Neuroscience*, vol. 2016, pp. 1–11, 2016.
- [12] M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep learning for tomato diseases: Classification and symptoms visualization," *Applied Artificial Intelligence*, vol. 31, no. 4, pp. 299–315, 2017.
- [13] H. Atila, M. Uçar, K. Akyol, and E. Uçar, "Plant leaf disease classification using EfficientNet deep learning model," *Ecological Informatics*, vol. 61, p. 101182, Jan. 2021.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1097–1105.
- [15] J. Deng et al., "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.
- [16] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3431–3440.
- [17] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 12, pp. 2481–2495, Dec. 2017.
- [18] G. Milioto, P. Lottes, and C. Stachniss, "Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in CNNs," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2229–22