



YOLOv8 assisted smart agriculture approach for tea (*Camellia sinensis*) leaf disease detection

SANJEET KUMAR BORAH^{1,2*}, KANDARPA KALITA¹, LAXMI NARAYAN SETHI² and AVINASH KUMAR²

Assam Down Town University, Panikhaiti, Guwahati, Assam 781 026, India

Received: 11 December 2025; Accepted: 03 June 2026

ABSTRACT

Tea (*Camellia sinensis* L.) is a major agricultural crop whose yield and quality are highly affected by foliar diseases that impair plant health and reduce market value. Many diseases exhibit similar early symptoms, making manual field inspection slow, inconsistent and prone to misdiagnosis under variable field conditions. These limitations create a strong need for accurate, scalable and automated disease detection methods to support effective tea crop management. To overcome these challenges, the present study was carried out during July–October 2025, that presents a YOLOv8-based deep learning framework capable of automatically detecting and localising multiple tea leaf diseases using images captured under natural field conditions. The dataset included seven major disease classes and a healthy leaf category with total 5,276 images, each annotated using normalised YOLO bounding-box formats and enhanced through diverse augmentation strategies to strengthen model generalisation. The YOLOv8-s model was trained for 100 epochs using a composite loss function that combines Complete IoU (CIoU) for accurate bounding-box regression with classification and objectness losses. Model performance was assessed using standard detection metrics. The trained model achieved a mean average precision @0.5 of 97.68%, a mean average Precision @0.5:0.95 of 97.68%, precision of 93.61% and recall of 95.31%, reflecting substantial improvement over the training period. Further diagnostic evaluations including confusion matrices, precision-recall curves and confidence-based performance analyses confirmed the model's robustness, even for visually similar disease symptoms such as brown blight and tea algal leaf spot. Overall, the proposed YOLOv8 framework demonstrated high accuracy and computational efficiency for tea leaf disease detection. Its strong performance under real field variability highlights its potential for integration into automated plantation monitoring systems, UAV-assisted crop surveillance and mobile-based advisory tools to support precision agriculture.

Keywords: Computer vision techniques, Deep learning models, Disease detection, Precision agriculture, YOLOv8 architecture

Tea (*Camellia sinensis* L.) is a plantation crop of global importance and a major stakeholder in India's agricultural economy (Pandey *et al.* 2021). Common tea leaf diseases like blister blight, red rust grey blight, brown blight, helopeltis and several other fungal and algal infections etiologically result in reduction of yield, decreased quality of leaf and increase the cost of management (Panthee *et al.* 2024). Conventional disease diagnosis is based on expert visual inspection, thus being labourious, subjective and not conducive to efficient monitoring of large plantation areas (Abbas *et al.* 2023). The precision agriculture is being widely used, therefore intelligent, automatic and high-precision disease detection technologies are the need of the hour for prompt action (Upadhyay *et al.* 2025). Advances

in artificial intelligence (AI), particularly deep learning, have revolutionised plant disease diagnosis by automating identification characteristics and dynamically classifying complex field situations. Similar successful results with Convolutional Neural Networks (CNN) have been reported for horticultural crops where CNNs achieved better results through handcrafted feature-based approach with high classification accuracy even in presence of lighting, occlusion and background noise changes (Vilar *et al.* 2024).

Recent applications of deep learning in leaf classification have been shown to be able to pick up slight, noticeable visual differences and disease symptoms over varied leaf datasets with potential for a global generalisation and robustness for agricultural applications (Srilatha and Praveen 2025). The innovation of object detection models like You Only Look Once (YOLO) family, especially the YOLOv8 has opened a new horizon for agricultural disease monitoring through accelerating real-time inference

¹Assam Down Town University, Guwahati, Assam; ²Assam University, Silchar, Assam. *Corresponding author mail: sanjeetkumarborah27@gmail.com

and precise mapping towards objects (Preanto *et al.* 2024). Many researchers like Wang *et al.* (2024) proposed a model that consists of a backbone and neck which work together to detect small and irregular tea leaf lesions with better accuracy. Its anchor free design also plays an important role for feature aggregation to support the detection process. The system runs fast and this makes it suitable for real-time field use. However, only a small number of studies have explored YOLOv8 for this purpose.

In this study, we have taken the advantages of a sophisticated detection on-the-fly using YOLOv8 model permitting faster inference and greater spatial resolution while being better-generalised than all existing types of CNN classifiers. The primary objective of this study was to utilise a deep learning-based object detection method for the precise identification of tea leaf illnesses from images captured under real field condition. This method improves the early block detection, automates on field monitoring system and also contributes a lot in data driven decision making in tea cultivation. The implementation of YOLOv8 in tea disease monitoring is a fast, scalable and intelligent system that fits well with the requirements of contemporary precision agriculture and sustainable plantation management.

MATERIALS AND METHODS

Experimental site: The present study was carried out during July–October 2025 at Assam Down Town University, Guwahati (26.2016°N, 91.8614°E; at an elevation of 59 m amsl), Assam. The study intends to evaluate the efficacy of the YOLOv8 model in identifying and categorising various disease classifications from the given dataset which was taken from Kaggle repository. The methodology of this work adheres to a methodical sequence comprising dataset description, pre-processing, augmentation, implementation, model training, performance evaluation and result analysis. Every stage was meticulously crafted to guarantee dependable detection and classification of tea leaf diseases in authentic field situations.

Dataset description: In this study, TeaLeafBD, an open source dataset from Kaggle repository, was used which contains class-wise collections of total 5,276 images of tea leaf which includes healthy images (935 images) and disease images like tea algal leaf spot (418 images), brown blight (506 images), grey blight (1,013 images), helopeltis (607 images), red spider (515 images) and green mirid bug (1,282 images), captured under real field conditions. In this study, the used dataset did not provide region-level annotations, which indicates where the disease symptoms appear in the




Fig. 1 Tea leaf images with different disease condition.

image. However, YOLOv8 model architecture is inherently an object detection model that requires bounding box annotations in the YOLO format. To overcome this issue, synthetic bounding boxes covering the entire image was adapted for YOLOv8 by assigning each image a synthetic bounding box that spans the entire frame. This allows the originally classification-focused dataset to be processed within a detection framework without altering the meaning of the visual content. All class folders are automatically read and indexed, and the resulting labels were stored in a dataset.yaml file to maintain reproducibility. The dataset was then divided into training, validation and testing subsets using a stratified 70-20-10 split, ensuring balanced representation across classes and reducing sampling bias. This setup creates a clean and scalable foundation for conducting YOLOv8-based experiments. Sample images for all tea leaf classes used in this study are presented in Fig. 1.

Pre-processing: A systematic pre-processing workflow was applied to prepare the data for efficient GPU training. Images of various formats were reorganised into YOLO-compatible directory structures, renamed using class-prefixed identifiers for traceability and resized to 640 pixels × 640 pixels using YOLOv8's letterbox method to preserve aspect ratio. Pixel values are normalised to the range of 0,1 to improve numerical stability during optimisation. Automatic Mixed Precision (AMP) was used to speed up training and reduce memory usage without affecting model accuracy. To ensure that results remain reproducible, fixed random seeds and controlled randomness were employed throughout the pre-processing and data-splitting stages.

Augmentation: To improve the model's ability to generalise diverse field conditions, YOLOv8 performs a set of controlled on-the-fly augmentations exclusively

on the training data. HSV-based adjustments introduce natural variations in colour and brightness, while spatial transformations such as horizontal flips, small translations and scaling help the model handle changes in viewpoint and leaf positioning. More aggressive operations including rotation, shear, perspective distortion and mix-up were intentionally disabled so that the disease-related texture patterns remain biologically realistic. Mosaic augmentation was used in the early training phase to blend four images into a single composite sample, expanding spatial diversity before it is gradually phased out in later epochs to stabilise learning. Altogether, this augmentation strategy strengthens the model's generalisation capability while preserving the integrity of disease features.

Implementation: In this study, a deep-learning based methodology built around the YOLOv8 object detection model was used for automatic detection of tea leaf diseases. The process began with collection of images from an existing dataset, where images were captured under natural lighting, varying leaf orientations, occlusions and diverse background environments. This variability ensures that the dataset represents realistic agricultural settings, allowing the model to generalise effectively under field conditions.

Each annotated instance within an image was represented using the YOLO bounding-box format. For an object belonging to class c , its annotation vector is expressed as:

$$A = \left(c, \frac{x}{W}, \frac{y}{H}, \frac{w}{W}, \frac{h}{H} \right)$$

Where (x,y) , Center of the bounding box; (w,h) Width and height; (W,H) , Corresponding image dimensions.

Normalisation ensures scale invariance, allowing the detector to operate consistently across images of different resolutions. A detailed inspection of bounding-box statistics confirmed broad variation in object size and location, supporting strong spatial generalisation.

To enhance representational richness and prevent over-fitting, the dataset was expanded through several augmentation strategies, including random rotations, contrast and brightness shifts, HSV adjustments and YOLOv8's built-in mosaic augmentation. These transformations effectively increased the underlying distribution $p(x|c)$, enabling the model to better approximate the posterior probability $p(c|x)$ in diverse field environments. The dataset was partitioned into training (70%), validation (20%) and testing (10%) subsets.

YOLOv8-s was selected for its computational efficiency and robustness in small-object detection tasks. The model adopts an anchor-free detection mechanism and uses C2f blocks in its backbone for enhanced gradient flow, while its PAN-FPN structure enables multi-scale feature fusion. Training was formulated as the minimisation of a composite loss function,

$$L = \lambda_1 L_{\text{box}} + \lambda_2 L_{\text{cls}} + \lambda_3 L_{\text{obj}}$$

Where L_{box} , Bounding-box regression loss; L_{cls} , Classification loss; L_{obj} , Objectness loss.

YOLOv8 uses the Complete IoU (CIoU) loss for bounding-box regression, expressed as:

$$L_{\text{CIoU}} = 1 - \text{IoU} + \frac{(\rho^2(b, b^*))}{c^2} + \alpha v$$

Where IoU, Intersection-over-Union between predicted and ground-truth boxes; ρ , Euclidean distance between their centers; c , Diagonal of the minimum enclosing box; v , An aspect-ratio consistency term modulated by:

$$\alpha = \frac{v}{(1 - \text{IoU}) + v}$$

Training was performed for 100 epochs using Stochastic Gradient Descent (SGD) with a learning rate of 0.01 and a batch size of 16. The model parameters were updated using the following optimisation rule:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} L(\theta_t)$$

Where θ_t , Model parameters at iteration t ; η , Learning rate; $\nabla_{\theta} L(\theta_t)$, Gradient of the loss function with respect to the model parameters. The negative sign indicates that the parameters are updated in the direction that minimises the loss.

Input images were resized to 640 pixels $\times$ 640 pixels to maintain consistency during training. Training and validation losses, along with precision (P), recall (R) and mean average precision (mAP) were monitored throughout training to assess model convergence and performance.

While the F1-score, which balances both, was defined as:

$$F_1 = \frac{2PR}{P + R}$$

Localisation accuracy was assessed using IoU, computed as:

$$\text{IoU} = \frac{|B_p \cap B_g|}{|B_p \cup B_g|}$$

where B_p and B_g , Predicted and ground-truth bounding boxes, respectively.

The Average Precision (AP) for each class was computed as the integral of the precision-recall curve,

$$\text{AP} = \int_0^1 P(R) dR$$

The overall mAP for N classes was obtained through,

$$\text{mAP} = \frac{1}{N} \sum_{i=1}^N \text{AP}_i$$

A stricter measure, $\text{mAP}@0.5:0.95$, averaged AP across IoU thresholds from 0.5 to 0.95.

To compare the method with existing approaches, performance of several widely used CNN-based classification models was compared with the proposed YOLOv8 model. Because CNNs work purely as image-level classifiers, their results are typically reported in terms of accuracy. YOLOv8, on the other hand, follows an object-detection framework, so its performance is measured using $\text{mAP}@0.5$, which accounts for both the predicted class and the confidence of the detection. Although accuracy and mAP represent




Fig. 2 Distribution of annotated samples across all disease categories and healthy leaves.

different evaluation styles, presenting them together helps clarify how these methods differ while keeping each model's original metric intact.

RESULTS AND DISCUSSION

The distribution of annotated samples across the dataset (Fig. 2), indicated that green mirid bug had the highest frequency, whereas tea algal leaf spot appeared the least. This inherent imbalance reflects natural field prevalence and explains the comparatively lower recall and average precision for minority classes (Table 1). Despite this, the dataset remained sufficiently diverse to support balanced training.

The normalised confusion matrix illustrated in Fig. 3 showed a strong diagonal pattern, confirming high class-wise accuracy. Healthy leaf achieved the highest accuracy at 99%, followed by helopeltis and red spider with accuracies of

Table 1 Class basis distribution of dataset

Class	No. of Images
Tea algal leaf spot	418
Brown blight	506
Gray blight	1,013
Helopeltis	607
Red spider	515
Green mirid bug	1,282
Healthy leaf	935
Total	5,276

94% and 92%, respectively. In previous studies related to leaf disease detection, it was found to be less accurate (Ruby *et al.* 2024). The limited misclassification observed occurred predominantly among visually similar categories, such as between tea algal leaf spot and background, and between brown blight and other foliar lesions which has also been reported by Ghafar *et al.* (2024). These patterns highlight the natural morphological overlap among disease symptoms.

Fig. 4 illustrates how the proposed model responded to changes in the confidence threshold, offering insight into its practical reliability beyond a single accuracy measure. As shown in Fig.

4(A), the macro-averaged F1-score remained, consistently high across a broad range of thresholds. This behaviour indicates that the model achieved a stable balance between precision and recall and was not sensitive to a specific threshold choice as reported by Berger and Guda (2020). The class-wise F1 curves followed a similar trend, suggesting that performance is reasonably uniform across disease categories, with no strong bias toward a particular class. Small variations at very low or very high confidence levels are expected and are typically associated with visually ambiguous samples.

The precision-confidence curves in Fig. 4(B) showed a clear and steady improvement in precision as the confidence threshold increases. This trend confirmed that predictions made with higher confidence are more reliable and less affected by false positives. From an application perspective, this is particularly relevant in agricultural disease monitoring, where incorrect disease alerts can lead to unnecessary spraying or management actions. Similar behaviour has also been observed in recent plant disease detection studies that emphasise confidence-aware evaluation for field deployment.

Fig. 4(C) presents the precision-recall curves used to compute AP and mAP values. The smooth shape of the curves and the absence of abrupt drops indicated stable class discrimination across different recall levels. This suggests that the model was able to preserve precision even when a larger number of diseased samples are retrieved. Compared with earlier CNN-based methods, which often show a sharp loss of precision at higher recall, the proposed approach demonstrated better control over this trade-off, especially under visually complex conditions.




Fig. 3 F1-confidence curve showing macro-averaged and class-wise F1-score variations across confidence thresholds.

The recall-confidence curves in Fig. 4(D) showed a gradual decline in recall as the confidence threshold increased, which is a normal and expected pattern. Notably, recall remained reasonably high at moderate confidence levels, implying that most diseased instances can still be detected without lowering the threshold excessively. This behaviour is advantageous for real-world use, where a single fixed threshold is often selected for simplicity and consistency.

Taken together, the curves in Fig. 4 indicate that the proposed model delivered stable and well-balanced performance across confidence thresholds. The consistency observed across F1, precision, recall and PR curves supports the robustness of the model and highlights its suitability for practical crop disease detection tasks under field conditions.

The evolution of performance metrics during training (Fig. 5),




Fig. 4 Model performance curves across confidence thresholds: (a) F1-confidence, (b) precision-confidence, (c) precision-recall (AP and mAP), and (d) recall-confidence curves.




Fig. 5 Evolution of core performance metrics across training epochs.

highlights monotonic improvements in precision, recall, mAP @0.5 and mAP @0.5:0.95. The close alignment of mAP @0.5:0.95 with mAP @0.5 further demonstrates that the model maintained strong localisation accuracy even under stricter IoU thresholds.

A comprehensive summary of model improvements

Table 2 Final training performance summary illustrating overall metric improvements and training duration

Metric	Initial value	Final value	Improvement
mAP @0.5	76.24%	97.68%	+21.45%
mAP @0.5:0.95	73.75%	97.68%	+23.93%
Precision	66.43%	93.61%	+27.18%
Recall	76.89%	95.31%	+18.42%
Total training time		5337 sec	~1.48 h

mAP, Mean average precision.

Table 3 Performance comparison of existing CNN-based classifiers with the proposed YOLOv8 detection approach

Sl. no.	Reference	Model used	Evaluation matrix	Accuracy
1.	Rahman <i>et al.</i> (2024)	CNN	Accuracy	96.70%
2.	Li <i>et al.</i> (2024)	Improved MobileNetV3	Accuracy	94.45%
3.	Sun <i>et al.</i> (2018)	CNN	Accuracy	93.75%
4.	Heng <i>et al.</i> (2024)	CNN	Accuracy	92.47%
5.	Datta and Gupta (2023)	CNN	Accuracy	96.56%
6.	Thangaraj <i>et al.</i> (2023)	M-RN50V2	Accuracy	90.84%
7.	Ayon <i>et al.</i> (2024)	ViT	Accuracy	96.09
8.	Srivastav <i>et al.</i> (2023)	CNN	Accuracy	84%
9.	Singh <i>et al.</i> (2023)	CNN	Accuracy	84.5%
10.	Zhang <i>et al.</i> (2023)	IEM-ViT	Accuracy	93.78%
11.	Proposed model	YOLOv8	Accuracy	97.68%

CNN, Convolutional Neural Network; ViT, Vision Transformer; IEM-ViT, Information Entropy Masked Vision Transformer; YOLO, You Only Look Once version 8.

(Table 2), showed substantial gains across all major performance indicators. Compared to initial values, mAP @0.5 increased by 21.45%, mAP @0.5:0.95 increased by 23.93%, precision improved by 27.18% and recall improved by 18.42% which is better than the results found in previous study by Olorunshola *et al.* (2023). The total training duration of approximately 1.48 h underscored the computational efficiency of YOLOv8 for large-scale agricultural applications.

Collectively, these findings confirmed that YOLOv8 delivers outstanding performance for automated tea leaf disease detection. Its strong classification accuracy, robust localisation capability and stable

training convergence established it as a powerful tool for deployment in UAV-based field monitoring, automated plantation assessment and mobile advisory systems supporting precision agriculture.

Performance comparison with existing work: The proposed YOLOv8 approach reached 97.68% mAP @0.5, performing better than the CNN models within their respective evaluation setups and showing strong capability in identifying tea leaf diseases (Table 3).

The experimental results indicated that the YOLOv8 model is capable of recognising tea leaf diseases with a level of consistency that is suitable for field application. The model performed reliably across most disease categories and maintained stable accuracy even under varying lighting conditions and complex background scenes. This robustness is particularly important for tea plantations, where early disease symptoms are often subtle and difficult to identify through visual inspection alone. Moreover, this approach

can improve overall crop yield and quality by increasing monitoring efficiency and timely decision making. The results suggest that the proposed approach can support timely disease monitoring in real plantation environments by introducing advanced systems like UAV (Unmanned aerial vehicle), and agricultural robots.

At the same time, the study highlights areas that require further improvement. Certain disease classes were represented by fewer training samples and their symptoms were visually faint, leading to occasional misclassification. Overlapping leaves and background blending also reduced detection clarity in some cases. Expanding the dataset with more representative images, especially from early infection stages, would help address these issues. Future work may also explore lighter models for mobile deployment, adaptation to seasonal variations and interpretability tools to enhance user trust and long-term reliability.

REFERENCES

- Abbas A, Zhang Z, Zheng H, Alami M M and Alrefaei A F. 2023. Drones in plant disease assessment, efficient monitoring and detection: A way forward to smart agriculture. *Agronomy* **13**(6): 1524. <https://doi.org/10.3390/agronomy13061524>
- Ayon R, Mubin A S, Hridoy Al-A G, Mojumder M U, Chakraborty N and Islam Md J. 2024. Deep learning for tea leaf disease detection: Comparing VGG16, vision transformer, efficientnetb3 and xception. (In) *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, Kamand, 24–28 June, pp. 1–6. <https://doi.org/10.1109/ICCCNT61001.2024.10726203>
- Berger A and Guda S. 2020. Threshold optimisation for F measure of macro-averaged precision and recall. *Pattern Recognition* **102**: 107250. <https://doi.org/10.1016/j.patcog.2020.107250>
- Datta S and Gupta N. 2023. A novel approach for the detection of tea leaf disease using deep neural network. *Procedia Computer Science* **218**: 2273–86. <https://doi.org/10.1016/j.procs.2023.01.203>
- Ghafar A, Chen C, Atif A S, Ur Rehman Z and Rahman G. 2024. Visualising plant disease distribution and evaluating model performance for deep learning classification with YOLOv8. *Pathogens* **13**(12): 1032. <https://doi.org/10.3390/pathogens13121032>
- Heng Q, Yu S and Zhang Y. 2024. A new AI-based approach for automatic identification of tea leaf disease using deep neural network based on hybrid pooling. *Heliyon* **10**(5): e26465. <https://doi.org/10.1016/j.heliyon.2024.e26465>
- Li Y, Lu Y, Liu H, Bai J, Yang C, Yuan H, Li X and Xiao Q. 2024. Tea leaf disease and insect identification based on improved MobileNetV3. *Frontiers in Plant Science* **15**: 1459292. <https://doi.org/10.3389/fpls.2024.1459292>
- Olorunshola O E, Irhebhude M E and Ewiewkpaefe A E. 2023. A comparative study of YOLOv5 and YOLOv7 object detection algorithms. *Journal of Computing and Social Informatics* **2**(1): 1–12. <https://doi.org/10.33736/jcsi.5070.2023>
- Pandey A K, Sinniah G D, Babu A and Tanti A. 2021. How the global tea industry copes with fungal diseases – Challenges and opportunities. *Plant Disease* **105**(7): 1868–79. <https://doi.org/10.1094/PDIS-09-20-1945-FE>
- Panthee D R, Pandey A and Paudel R. 2024. Multiple foliar fungal disease management in tomatoes: A comprehensive approach. *International Journal of Plant Biology* **15**(1): 69–93. <https://doi.org/10.3390/ijpb15010007>
- Preanto S A, Ahad M T, Emon Y R, Mustofa S and Alamin Md. 2024. A semantic segmentation approach on sweet orange leaf diseases detection utilising YOLO. *arXiv*. <https://doi.org/10.48550/arXiv.2409.06671>
- Rahman H, Ahmed J, John P, Salam M A and Rabbi M F. 2024. Automated detection of selected tea leaf diseases in Bangladesh with convolutional neural network. *Scientific Reports* **14**: 14097. <https://doi.org/10.1038/s41598-024-62058-3>
- Ruby E D K, Amrithayogam G, Sasi G, Chitra G, Choubey A and Gopalkrishnan S. 2024. Advanced image processing techniques for automated detection of healthy and infected leaves in agricultural systems. *Mesopotamian Journal of Computer Science* **4**(1): 44–52. <https://doi.org/10.58496/MJCSC/2024/006>
- Singh R, Sharma N, Chauhan R, Choudhary A and Gupta R. 2023. Automatic approach based on deep learning for tea leaf disease detection. (In) *2nd International Conference on Futuristic Technologies (INCOFT)*, Belagavi, Karnataka, 24–26th November 2023, pp. 1–5. <https://doi.org/10.1109/INCOFT60753.2023.10425660>
- Srilatha A and Praveen P. 2025. CNN-based plant leaf disease identification and classification for sustainable agriculture. (In) *6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI)*, Goathgaun, Nepal, 7–8th January 2025, pp. 1459–65. <https://doi.org/10.1109/ICMCSI64620.2025.10883522>
- Srivastav S, Guleria K and Sharma S. 2023. Tea leaf disease detection using deep learning-based convolutional neural networks. (In) *2023 IEEE World Conference on Applied Intelligence and Computing (AIC)*, Varanasi, 29–30 July, pp. 569–74. <https://doi.org/10.1109/AIC57670.2023.10263835>
- Sun X, Mu S, Xu Y, Cao Z and Su T. 2018. Image recognition of tea leaf diseases based on convolutional neural network. (In) *International Conference on Security, Pattern Analysis and Cybernetics (SPAC)*, Jinan, China, 14–17 December, pp. 309. <https://doi.org/10.1109/SPAC46244.2018.8965555>
- Thangaraj R, Vadivel B, Arvind T, Chandhru L, Harish K and Eswaran S. 2023. Tea leaf disease identification using improved convolution neural network. (In) *International Conference on Digital Applications, Transformation and Economy (ICDATE)*, Miri, Malaysia, 14–16 July, pp. 1–5. <https://doi.org/10.1109/ICDATE58146.2023.10248508>
- Upadhyay A, Chandel N S, Singh K P, Chakraborty S K and Nandede B M. 2025. Deep learning and computer vision in plant disease detection: A comprehensive review of techniques, models and trends in precision agriculture. *Artificial Intelligence Review* **58**(3): 92. <https://doi.org/10.1007/s10462-024-11100-x>
- Vilar A M, Garcia L, Garcia S A, Asorey C R and Garcia H J. 2024. Enhancing precision agriculture pest control: A generalised deep learning approach with YOLOv8-based insect detection. *IEEE Access* **12**: 84420–34. <https://doi.org/10.1109/ACCESS.2024.3413979>
- Wang C, Li H, Deng X, Liu Y and Wu T. 2024. Improved you only look once v.8 model based on deep learning: Precision detection and recognition of fresh leaves from yunnan large-leaf tea tree. *Agriculture* **14**(12): 2324. <https://doi.org/10.3390/agriculture14122324>
- Zhang J, Guo H, Guo J and Zhang J. 2023. An information entropy masked vision transformer (IEM-ViT) model for recognition of tea diseases. *Agronomy* **13**(4): 1156. <https://doi.org/10.3390/agronomy13041156>