

Artificial Intelligence for Crop Pest Detection: A Review of Deep Learning Techniques, Challenges, and Future Directions

Deshmukh Nishant Dhananjay, Prof. Abhishek Popat Bangar, Prof. S. B. Bhosale, Prof. Karad V. A.

Department of Computer Engineering
Jaihind College of Engineering, Kuran, Maharashtra, India

Abstract — Pests reduce both the yield and the quality of crops, and checking fields by hand is slow, tiring, and dependent on trained experts. Artificial intelligence (AI), particularly computer vision and deep learning, makes it possible to recognize pests from photographs and to respond to infestations earlier and with better precision. This review looks at AI-driven pest detection in terms of image classification, object detection, public benchmark datasets, model efficiency, and practical deployment. A focused narrative search of scholarly sources and publisher pages was carried out, covering work on the IP102 benchmark, the Pest24 dataset, machine-learning and deep-learning surveys, YOLO-based pest detectors, and rice-pest detection. The literature shows that convolutional neural networks can learn the visual cues needed for pest classification, while object detectors from the YOLO family can find and label several pests within one image. The choice of dataset strongly shapes how trustworthy an evaluation is: IP102 holds more than 75,000 images in 102 categories, with bounding boxes for roughly 19,000 of them, whereas Pest24 consists of field images in which pests are small and often packed closely together. Newer model-specific studies report gains on the test sets they selected, but figures drawn from different datasets and evaluation protocols cannot be compared head to head. Challenges that keep recurring include tiny pests, look-alike species, imbalanced classes, cluttered field backgrounds, changing light, a shortage of labelled field data, and the limited computing power available on devices. The review points to the need for more representative field testing, open evaluation practices, efficient models, reporting of uncertainty, and stronger links between detection output and integrated pest management. AI-based detection is promising, yet a high benchmark score alone does not prove dependable performance on real farms. Future work should focus on external validation, ease of use for farmers, and responsible use of model predictions.

Keywords— crop pest detection, artificial intelligence, deep learning, computer vision, YOLO, IP102, precision agriculture

I. INTRODUCTION

Pests are a constant threat to farming, since they can harm leaves, stems, roots, flowers, and fruit. Spotting them early helps growers judge whether a closer look or some form of control is warranted. Monitoring has traditionally depended on walking the field, setting traps, and relying on experts to inspect plants by eye. These practices are still valuable, but they take a great deal of time and are hard to extend over wide areas or to repeat as often as monitoring demands [1,2].

AI-based pest detection usually works with images from smartphones, fixed cameras, insect traps, drones, or other monitoring devices. The model takes an image and returns a pest category, the position of each insect it finds, or both. Classification tells us which pest appears in a picture; object detection also estimates where each target sits; and segmentation marks the exact pixels that belong to it. Because these tasks serve different purposes, they should not be judged by the same yardstick.

Deep learning has become popular because convolutional neural networks (CNNs) can pick up visual patterns directly from labelled images.

Detectors such as the You Only Look Once (YOLO) family are appealing when both speed and localization matter. Farm images, however, are unlike the pictures found in common benchmarks: insects can be minute, partly hidden, tightly clustered, or camouflaged against plants and soil.

A system that works well on one dataset may be less dependable on another crop, location, camera, or season [2–5]. This review aims to summarize representative AI approaches to crop pest detection, compare selected studies and datasets, point out recurring gaps, and suggest practical directions for future work. It concentrates on image-based pest identification and detection; it does not propose or experimentally test a new detection model.

II. REVIEW METHODOLOGY AND LITERATURE SELECTION CRITERIA

This paper follows a focused narrative review approach rather than a registered systematic review or meta-analysis. Relevant records were located through targeted searches of scholarly search results, publisher pages, open-access journal pages, and conference repositories, using combinations of the phrases “crop pest detection,” “deep learning,” “machine learning,” “YOLO,” “IP102,” “Pest24,” and “rice pest detection.” Sources were chosen to cover three groups: (i) broad reviews, (ii) public benchmark datasets, and (iii) model-specific detection studies. A study was treated as relevant if it dealt directly with image-based recognition or detection of insect pests in agricultural settings, introduced a dataset or model, or reviewed methods and their limitations. Papers about unrelated computer-vision tasks were left out. Preference went to sources with accessible abstracts or full text and clear publication details. The search is a targeted starting point, not an exhaustive one: no formal duplicate screening, independent reviewer process, or statistical meta-analysis was carried out. The findings should therefore be read as a qualitative synthesis of the selected literature. Before academic submission, the search should be broadened and the final list of papers checked against the supervisor’s requirements.

III. LITERATURE REVIEW

1. Machine learning and deep learning approaches

Early image-recognition pipelines relied on hand-crafted features such as color, texture, shape, or local descriptors, which were then passed to a classifier like a support vector machine or a decision tree. These methods can work well when imaging conditions are controlled, but their results depend heavily on how the features are designed and on image quality. Deep-learning methods instead learn features straight from training examples and can capture far more complex visual patterns. Chithambarathanu and Jeyakumar [1] survey machine-learning and deep-learning approaches to identifying crop pests and pathogens, covering CNN-based methods as well as conventional classifiers. Their survey shows how wide the range of techniques tried across different crops and pest categories has become.

2. Datasets and benchmark design

A model’s results depend not only on its architecture but also on the data used to train and test it. Wu et al. [2] introduced IP102, a benchmark of more than 75,000 images spread over 102 insect-pest categories, with bounding boxes annotated on about 19,000 of them. Its class distribution is naturally

imbalanced, and the authors single out variation between and within classes as key difficulties. IP102 supports both fine-grained pest recognition and object-detection research, yet no benchmark automatically mirrors every field environment or local pest population.

Wang et al. [3] presented Pest24, a multi-target dataset built from field pest-monitoring imagery. It covers 24 pest categories and stresses very small objects, dense groupings, and look-alike pests. The authors compared several detection methods, including Faster R-CNN, SSD, YOLOv3, and Cascade R-CNN. The dataset matters because it exposes difficulties that may be hidden in collections where a single large, clearly visible insect fills the frame.

3. YOLO-based and lightweight detection models

YOLO-based detectors are widely studied for real-time pest detection because they predict object locations and classes in one pass. Modifications usually aim at better feature extraction, multi-scale feature fusion, attention mechanisms, or fewer parameters. A YOLOv5-based study, “Pests Identification of IP102 by YOLOv5 Embedded with the Novel Lightweight Module,” explores a lightweight model for recognizing a broad range of pest categories [4]. Work of this kind tackles a practical tradeoff: a model has to be accurate enough to be useful, yet light enough to run on the hardware people actually have.

4. Pest detection for specific crops

Crop-specific studies can tune a model to the pests, backgrounds, and monitoring needs of one crop. Xiong et al. [6] examined rice-pest detection using a modified YOLOv8 model with attention and feature-fusion components. On a rice-pest subset of IP102 with 9,663 images and eight pest types, they reported an mAP@0.5 of 98.8% and an mAP@0.5:0.95 of 78.6%. The gap between these two figures is worth noting: mAP@0.5 applies a more lenient overlap threshold than mAP@0.5:0.95, so the two cannot be swapped for one another. The study illustrates task-specific model improvement, but independent testing would still be needed to establish how it performs on new farms and under different imagecapture conditions.

5. Broader reviews and emerging directions

Recent reviews look at AI for plant disease and pest recognition across classification, detection, segmentation, and other sensing methods. Shoaib et al. [7] summarize deep-learning approaches and discuss the practical obstacles to real-world use. A 2025 review, “Research Progress of Deep Learning-Based Artificial Intelligence Technology in Pest and Disease

Detection and Control,” also considers links with remote sensing, spectral imaging, and Internet of Things systems [8]. Together, these reviews suggest that pest detection is increasingly seen as one part of a larger monitoring pipeline and not as a stand-alone image-classification problem.

IV. COMPARATIVE ANALYSIS OF EXISTING STUDIES

Table 1 sets selected studies side by side according to their main contribution. Since the studies differ in datasets, pest categories, train/test splits, and metrics, their reported numbers should not be ranked as though they came from one controlled experiment.

Table 1. Comparison of Selected Studies

Study	Focus / Data	Approach	Reported contribution	Main comparison caution
Chithambarathanu & Jeyakumar (2023) [1]	Survey across crop pests and pathogens	Conventional ML and deep learning	Maps the major families of approaches	Broad scope; results do not come from one common benchmark
Wu et al. (2019) [2]	IP102; 102 categories, >75,000 images	Dataset and baseline experiments	Offers a major benchmark; ~19,000 boxes	Natural class imbalance and wide image variation
Wang et al. (2020) [3]	Pest24; field imagery, 24 categories	Faster R-CNN, SSD, YOLOv3, Cascade R-CNN	Targets tiny, dense, lookalike pests	Trap imagery from the field may differ from phone photos
Dong et al. (2024) [5]	Custom dataset and IP102 evaluation	Lightweight YOLOv5based PestLite	Reports gains in accuracy and parameter efficiency	Results depend on dataset and evaluation protocol
Xiong et al. (2024) [6]	9,663 images; eight ricepest types	Modified YOLOv8 with attention and feature fusion	Reports mAP@0.5 of 98.8% on its own setup	Crop-specific subset; external validation still matters

1. Overall comparison

Taken together, the studies show progress in both recognition quality and computational efficiency. Dataset papers supply shared evaluation resources, while model papers adapt existing architectures to particular conditions. Even so, one strong metric does not prove overall superiority. A fair comparison would use the same data split, annotation protocol, input resolution, hardware, and evaluation metrics, and where possible would also report false negatives, false positives, inference speed, and results for each pest class.

V. RESEARCH CHALLENGES AND IDENTIFIED GAPS

1. Small and densely grouped pests

Many insects take up only a handful of pixels, and several may overlap one another. Downsampling inside a model can wipe out useful detail. Pest24 deliberately highlights small object size, density, and objects sticking together [3]. Future systems need robust multi-scale feature extraction and should be tested on crowded images.

2. Class imbalance and visually similar species

Some pest categories have far more images than others, so a model may learn the common ones well while overlooking rare but important pests. IP102 documents a naturally long-tailed

distribution along with variation within and between categories [2]. Studies should report per-class recall, class-balanced sampling, and carefully designed augmentation, and not lean on overall accuracy alone.

3. Domain shift and limited field validation

Images gathered in laboratories, taken from the internet, or shot under controlled lighting may not match field conditions. Background, weather, camera quality, crop growth stage, pest life stage, and geography can all change the image distribution. Many published results rest on one chosen dataset split; testing

on independent farms, seasons, and devices is needed before any claim of general reliability.

4. Accuracy, speed, and device limitations

High-capacity models can demand more memory, power, and processing time than a low-cost phone or edge device offers. Lightweight models ease some deployment barriers, though compression may also lower accuracy. Alongside detection metrics, papers should state model size, memory use, latency, and the hardware used.

5. Evaluation consistency and practical usefulness

Results are hard to compare when papers differ in pest categories, datasets, splits, image resolutions, and metrics. $mAP@0.5$ is not the same as $mAP@0.5:0.95$, and classification accuracy cannot be set directly against object-detection mAP . Moreover, spotting a pest does not by itself reveal how severe an infestation is or what treatment is best. Practical systems should show their confidence, support verification by an expert, and avoid presenting uncertain predictions as firm advice.

VI. PROPOSED RESEARCH DIRECTIONS

Drawing on the reviewed studies and the gaps above, the following directions look useful for future work:

1. Priorities

Field-first datasets: Collect images across different farms, crop stages, seasons, cameras, and lighting conditions, with documented annotation quality and class distributions.

External validation: Test models on farms or locations that were not part of training, and publish in-domain and out-of-domain results separately.

Small-object methods: Explore multi-scale features, higher-resolution crops, tiling, and detection heads built for small targets, and measure whether the added computation pays off.

Balanced and low-data learning: Look into class-aware sampling, few-shot learning, transfer learning, and carefully checked synthetic augmentation for rare pest categories.

Efficient deployment: Compare compact architectures, quantization, pruning, and edge inference on realistic hardware, reporting latency, power, memory, and accuracy together.

Trustworthy outputs: Calibrate confidence scores, flag uncertain or unfamiliar inputs, and offer visual explanations that help users inspect detections.

Decision support: Link pest detection with field scouting and integrated pest management guidance, while leaving treatment decisions to local agronomic recommendations and human oversight.

A practical future system could pair a smartphone camera with a lightweight detector and a simple interface showing the predicted pest category, bounding box, confidence score, and a prompt to verify uncertain cases. Such a system should first be assessed as a detection aid and not as an autonomous pesticide-prescription tool. A pilot study should measure technical performance as well as usability with the people who will actually use it.

VII. CONCLUSION AND FUTURE SCOPE

This review examined AI-based crop pest detection, paying attention to algorithms, datasets, comparative evaluation, and deployment challenges. The selected literature indicates that CNN-based methods and object detectors such as YOLO can automate parts of pest recognition and localization. IP102 and Pest24 offer important evaluation resources, and recent model-specific studies show that architectural changes can raise performance in particular test settings. Still, reported results vary with dataset design and evaluation protocol, and strong benchmark performance does not guarantee dependable results in the field.

The main gaps involve tiny and overlapping insects, imbalanced categories, look-alike species, changing field conditions, little external validation, and the limits of low-cost devices. Future work should give priority to representative field data, consistent reporting, independent testing, efficient deployment, handling of uncertainty, and integration with responsible pest-management workflows. The long-term aim should be more than a model with a high accuracy score; it should be a dependable tool that helps growers notice possible pest problems sooner and make better-informed decisions. The present paper is a focused narrative review, and a submission-ready version should be strengthened with a broader, documented search and checked against the supervisor's final requirements.

REFERENCES

1. M. Chithambarathanu and M. K. Jeyakumar, "Survey on crop pest detection using deep learning and machine learning approaches," Multimedia Tools and Applications, 2023, doi: 10.1007/s11042-023-15221-3.

2. X. Wu, C. Zhan, Y.-K. Lai, M.-M. Cheng, and J. Yang, "IP102: A Large-Scale Benchmark Dataset for Insect Pest Recognition," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, doi: 10.1109/CVPR.2019.00899.
3. Q.-J. Wang et al., "Pest24: A large-scale very small object data set of agricultural pests for multi-target detection," Computers and Electronics in Agriculture, 2020, doi: 10.1016/j.compag.2020.105585.
4. "Pests Identification of IP102 by YOLOv5 Embedded with the Novel Lightweight Module," Agronomy, vol. 13, no. 6, Art. no. 1583, 2023, doi: 10.3390/agronomy13061583.
5. Q. Dong et al., "PestLite: A Novel YOLO-Based Deep Learning Technique for Crop Pest Detection," Agriculture, vol. 14, no. 2, Art. no. 228, 2024, doi: 10.3390/agriculture14020228.
6. P. Xiong et al., "Deep learning-based rice pest detection research," PLOS ONE, vol. 19, no. 11, e0313387, 2024, doi: 10.1371/journal.pone.0313387.
7. M. Shoaib et al., "Leveraging deep learning for plant disease and pest detection: a comprehensive review and future directions," Frontiers in Plant Science, vol. 16, 2025, doi: 10.3389/fpls.2025.1538163.
8. "Research Progress of Deep Learning-Based Artificial Intelligence Technology in Pest and Disease Detection and Control," Agriculture, vol. 15, no. 19, Art. no. 2077, 2025, doi: 10.3390/agriculture15192077.